Contents

How these notes are coming to exist

It is Fall 2007. QFT classes are being taught at UC Berkeley by Richard Borcherds (RB), Nicolai Reshetikhin (NR), and Peter Teichner (PT). Anton LaTeXs these notes in class and edits them later.[1] The version you're currently reading was compiled October 4, 2014. They should be available at
http://math.berkeley.edu/~anton/index.php?m1=writings.

– When something doesn't make sense to me, I mark it with three big, eye-catching stars [[★★★ like this]]. If you can clear any of these up for me, let me know.

– If you have notes that I'm missing or if you have a correct/clear explanation for something which is incorrect/unclear, let me know (either tell me what you'd like to modify, give me some notes to go on, or update the tex yourself and send me a copy). Real (mathematical) errors should be fixed because it would be immoral to let them propagate (er . . . that is, sit there), and typographical errors hardly take any time to fix, so you shouldn't be shy about telling me about them.

[1] With the exception of NR22, which was done by Chris Schommer-Pries.

1 RB 08-28

This won't really be a seminar, it will be a short course on quantum field theory for mathematicians. Aim of the course: give mathematical answers to the following questions

1. What is a QFT? There are many incompatible answers.
2. How do you construct them? Nobody knows how to construct them non-perturbatively, so we'll do everything perturbatively. What is a Feynman measure? It is easy to prove they don't exist.
3. What is renormalization and reregularization?
4. What is gauge-invariance? Anomalies.

In the first seminar, we'll try to give a quick survey without proofs.

Recall what a classical field theory is. There are two basic ingredients for classical field theory.

1. Fiber bundle.
2. a Lagrangian.

A fiber bundle is locally (on M) something of the form $F \times M \to M$. We call M the base space and F the fiber., so we have a copy of F sitting over every point in M.

A *classical field* is a section of the fiber bundle. In the case of a product, this is just a function $M \to F$.

Example 1.1. Classical mechanics is a (trivial) field theory. Take $M = \mathbb{R}$ (thought of as time) and F is configuration space, which is typically a finite-dimensional manifold (the possible positions of some mechanical system). The fields are maps from $\mathbb{R}$ (time) to the configuration space F. ◇

Example 1.2. Statistical field theory. Take M = space (say $\mathbb{R}^3$) and $F = \mathbb{R}$ (say). Then a field is a real function on $\mathbb{R}^3$ ◇

Example 1.3. Quantum field theory. Take M = spacetime (some Lorentzian manifold, usually flat Minkowski space $\mathbb{R}^{1,3}$). Unfortunately, one of the major unsolved problems in physics is whether it is $\mathbb{R}^{3,1}$ or $\mathbb{R}^{1,3}$. Take $F = \mathbb{R}$ for a Hermitian scalar field. In this case, a classical field is just a real function on $\mathbb{R}^{1,3}$. ◇

In more complicated classical field theories, you could take $F = SU(3)$ (which underlies quantum chromodynamics). Most of the problems of field theory show up in the simple case where you take $F = \mathbb{R}$.

Remark 1.4. More generally, instead of taking total space to be $F \times M \to M$, you take some twisted version of it (a more general fiber bundle $E \to M$, which locally looks like $F \times M \to M$). Physicists almost always take the fiber bundle to be a product. ◇

The next ingredient is a Lagrangian.

Example 1.5. Take statistical field theory. $L(\phi) = \sum_i \left(\frac{\partial \phi}{\partial x_i}\right)^2 + m^2\phi^2 + \lambda\phi^4$. Here ϕ is a field (e.g. a real-valued function). This Lagrangian is a function from L: Fields×Base space→ $\mathbb{R}$, but it isn't any old function; it is sort of local. L depends only on ϕ and its derivatives at the point $x \in M$. This means that L is a function on the *Jet space* of the fiber bundle. ◇

A jet space is just the set of pairs $(\phi, x) \in$ (functions$\times M$), where we identify (ϕ_1, x) with (ϕ_2, x) if ϕ_1 and ϕ_2 agree to all orders at the point x (i.e. all their derivatives agree). There are also finite order jet spaces where you only look at a finite number of derivatives (identify (ϕ_1, x) with (ϕ_2, x) if all derivatives to third order are the same). Finite jet spaces aren't as nice as infinite order jet spaces.

There are some variations on Lagrangians. We have Lagrangians, Lagrangian densities, and actions. The purpose of a Lagrangian is to produce an action, which is what gets you into business. A Lagrangian is a function on jet space. You want to integrate the Lagrangian over spacetime (the base space) to get an action. So the action of a classical field is "$\int_M L(\phi)$". However, you can't integrate functions: you often explain to undergraduates that things like $\int x^2$ don't make sense; you need to multiply a function by a form before you can integrate it. A Lagrangian

density is a form on the manifold times a Lagrangian. This is a bit untidy. We have Jet space$\to M$, so if we have a form on M, we can pull it back to Jet space, so you can think of a Lagrangian density as a special kind of form on Jet space.

Often, M has a canonical n-form, in which case you can identify Lagrangians with Lagrangian densities. For example, if M has a metric and an orientation, this gives a volume form. You do need to worry about the difference if you're studying gravity, because then M has no canonical metric (the point being that gravity varies with the metric on M). In that case, Lagrangians and Lagrangian densities are not the same thing, and the right thing to use is a Lagrangian density. But usually, you don't need to worry too much about the difference.

Finally, what is an action? If you have a Lagrangian density, you can integrate it over spacetime to get an action $\int_M L(\phi)$, except you can't because M is usually non-compact, so there is no reason for the integral to converge. Most physics books ignore the problem by pretending the integral converges even when it is clear it doesn't. What seems to be going on is that even though you can't define the action of a field, you can define the DIFFERENCE of the actions of two fields ϕ_1 and ϕ_2 provided that they differ on a compact set.

So if we've written down a Lagrangian density, we can define variation. Now you're in business. You can define the classical equations of motion (Euler-Lagrange equations): $\int L(\phi)$ is STATIONARY under variations of ϕ on COMPACT SETS.

For classical mechanics, you need a fiber bundle and a Lagrangian density. Once you have these things, you hand them to somebody who knows classical mechanics and they'll get excited and solve the equations of motion for you.

Now let's talk about what a quantum field theory is (let's do the case of a Hermitian scalar field $\phi : \mathbb{R} \times M \to M$). The basic idea: we should have (1) a Hilbert space H and (2) an operator $\phi(x)$ on H for each $x \in M$ satisfying various axioms. The problem is that it is impossible to make sense of the operator $\phi(x)$. The problem is that $\phi(x)\phi(y)$ has really dreadful singularities as x and y get close together. All of quantum field theory is in some sense trying to get around the problem of how to define $\phi(x)\phi(x)$. You get around this in two steps. The first thing you do whenever you have singularities is to smooth them out by convolving with smooth functions with compact support. Instead of using operators $\phi(x)$ for $x \in M$, we use $\phi(f) =$"$\int \phi(x) f(x)\, d^n x$" where f is smooth with compact support. There is another problem: even smoothed operators $\phi(f)$ are not defined on H. This turns out to be a fairly minor problem. They are defined on a dense subset $D \subseteq H$ and map D to D (they are "unbounded operators").

So what you end up with is an unbounded operator on H for each smooth compactly supported classical field f.

Now we'll give a minimal definition of a quantum field theory. We need (1) a module D over some $*$-commutative ring (in practice, you do perturbative stuff, so you use formal power series) (a $*$-ring is a ring with involution $*$ so that $(ab)^* = b^*a^*$), (2) a hermitian inner product $(\cdot,\cdot)\colon D \times D \to \mathbb{R}$, and (3) a $*$-algebra A of operators acting on D (generated by $\phi(f)$). Furthermore, $(Ax, y) = (x, A^*y)$.

This is the minimal amount of stuff you need to reasonably say you have a quantum field theory.

Example 1.6. QFT satisfying the Wightman axioms (we'll discuss these later). In this case, D is a dense subset of a Hilbert space and A is the algebra generated by the field operators. ◇

Example 1.7. Let L be a Lie algebra acting on a vector space D with an invariant symmetric inner product. Then you can construct an algebra by taking the universal enveloping algebra $U(L)$ with $a^* = a$ for all $a \in L$. ◇

Example 1.8. If G is a group and D is an orthogonal representation of G, then we can take A to be the group-ring of G with $g^* = g^{-1}$ for all $g \in G$. ◇

Example 1.9 (generalizing the last two examples)**.** A a Hopf algebra and D to be an orthogonal representation of A. ◇

Example 1.10. Take A to be any C^*-algebra or von Neumann algebra and D to be any Hilbert space that is a $*$-representation of A. ◇

Warning 1.11. People often define quantum field theories in terms of C^*-algebras. You have to watch out, because in our cases, the operators will usually be unbounded, and the C^*-algebra examples have bounded operators. ⌟

To construct a QFT, we need to give (1) a $*$-algebra A, and (2) a module D, and these things should satisfy some extra axioms.

How to construct examples. The algebra A is easy to construct. You could just take it to be the universal $*$-algebra generated by classical fields. The hard part is to construct the right representation. Start with a state w (a linear map $w\colon A \to R$ such that $w^* = w$) and define $(\cdot,\cdot)$ on A, given by $(a,b) = w(ab^*)$ and take $D = A/\ker(,)$ (this is basically the GNS construction). ω is constructed using Feynman integral, which is formally given by $\int e^{i\int L(\phi)} D\phi$ which can be expanded as a series of Feynman diagrams.

2 RB 09-04-2007

Today we'll continue with what should have been the second half of the first lecture.

Recall that last time we said that a classical field theory consists of (1) a fiber bundle over some manifold M (which might be space-time) and (2) a Lagrangian density, which is a special kind of form on Jet space.

For a quantum (or statistical) field theory, you have to specify a $*$-algebra A acting on some module with a sesquilinear form $(,)$ satisfying some axioms which we'll talk about next week. This algebra will usually be generated by classical fields.

Roughly, to get a form, you take a state $\omega\colon A \to R$ (with $R = \mathbb{R}$ or $\mathbb{R}[[\lambda]]$) on A such that $\omega^* = \omega$ ($\omega^*(a) = \omega(a)$). Then you can define $(a,b) = \omega(ab^*)$. So the main problem is to construct a state ω on the algebra A. The state ω is constructed (at least formally) in terms of Feynman integrals, which look something like

$$\int (\int \phi(x)^* f(x) dx) e^{i\int L(\phi) d^4x} D\phi$$

You integrate over the space of all fields. A problem: there is no invariant measure on infinite-dimensional spaces. To get around this, think about what a measure on a space X really is. It is a map from some subsets of X to $\mathbb{R}$ which is countably additive, etc. This approach to thinking about a measure doesn't work very often when you work on an infinite-dimensional space. A "Radon measure" is a linear map from continuous functions with compact support to $\mathbb{R}$, given by taking f to a real number which you think of as $\int f d\mu$. If you have a usual measure, then usual measure theory tells you how to construct the integral of a reasonable function. However, it turns out that you can't even define a Radon measure on an infinite-dimensional space (there probably aren't any continuous functions with compact support except for the zero function).

Instead: define a measure to be a linear map from *some space* of functions to $\mathbb{R}$. We only define the integral of functions we are interested in. But which functions are you interested in? Thinking about it, you see that most things you're interested in are of the form

$$\int (\int \phi^4(x) f(x) dx)(\int \phi^2(x) \partial\phi(x) dx) \cdots e^{i\int \text{quadratic}\, dx} D\phi$$

Where the first functions are given by integrating a form on Jet space. So we want to integrate polynomials in (certain forms on Jet space)$\times e^{\text{quadratic}}$. A *Feynman measure* is a linear functional on this space with certain properties (which we'll talk about later). This is a perfectly rigorous concept. You have to remember that physicists are doing perfectly reasonable mathematics, but they lie and say that they're doing something else. If you examine it carefully, a physicist means by "measure" a linear function on this space. PT: what if the fiber is not linear? RB: you want the fiber to at least be an affine space and you choose a vacuum to turn it into a vector space.

Next problem: is the Feynman measure unique? This is one of the most confusing things about quantum field theory. (1) Most physics books imply that it is, but they are well aware it is NOT. We need to explain why it behaves as if it is unique. Why can we get away with pretending it is unique? The key point is the following: There is an (infinite-dimensional) group of *renormalizations* which acts simply transitively on the Feynman measures. So there is a unique well-defined *orbit* of Feynman measures.

If you try to read a physics book, it gives the following misleading picture. If you start with a Lagrangian, you get a QFT (this is wrong). The correct picture: If you start with a Lagrangian and a Feynman measure, you get a QFT. Furthermore, there is a group acting on Lagrangians and Feynman measure which preserves the QFT, so you should take (Lagrangians)$\times$ (Feynman measures)/(Group of renormalizations). The group of renormalizations is the same size as the renormalization group, so you can often get away with just picking a Lagrangian. AJ: what you're calling "Feynman measure" is what is usually called a regularized path integral? RB: yes, a choice of Feynman measure is equivalent to a choice of regularization and renormalization scheme.

Toy example. Can we find an invariant measure on a 1-dimensional space? Yes, but there is no canonical way to do it if you want a translation-invariant measure: if $d\mu$ is a translation-invariant measure, so is $sd\mu$. Rescaling acts on the space of all translation-invariant measures. Furthermore, the group of rescalings acts simply transitively on the invariant measures on $\mathbb{R}^1$, so it doesn't really matter which measure you take most of the time. The same sort of thing happens in QFT, but the group of renormalizations is infinite-dimensional and non-abelian.

Consequences of non-uniqueness of Feynman measures are things called ANOMALIES. An anomaly is given by the following. Suppose the Lagrangian L is invariant under some group G (say G is the group of gauge transformations). We would expect G to act on the QFT of L. There is a problem because the QFT depends on the Feynman measure as well as L. Can we find a Feynman measure which is also invariant under G (and hence get a G-invariant QFT)? Sometimes you can and sometimes you can't. What is the obstruction to doing this?

We have a group H which acts simply transitively on a space X. In this case, by choosing a point of X, you can identify X with H. Now suppose another group G acts on H and X. Can we identify X with H in a G-invariant way? The answer is that there is an obstruction given by an element of the non-abelian cohomology group $H^1(G, H)$ (so you can do it when this element is zero in the cohomology group). This group H will be the group of renormalizations and G will be the group of symmetries. The problem of whether you can get G acting on the QFT amounts to working out an element of $H^1(G, H)$ and checking if it is zero. Thus, anomalies are elements of $H^1(G, H)$.

Now I'll try to explain quickly what a Feynman diagram is. A Feynman integral can be formally expanded as an infinite series of Feynman diagrams. A Feynman diagram looks like  . Each Feynman diagram is an abbreviation for a finite-dimensional integral. Each point represents a point of spacetime and each line represents a propagator $\Delta(x, y)$ (which is a Green's function for the quadratic part of the Lagrangian); just think of it as some function of x and y. Then the Feynman diagram 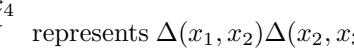

represents $\Delta(x_1, x_2)\Delta(x_2, x_3)^3\Delta(x_3, x_4)$ integrated over some of the x_i's. You usually know that this integral does NOT converge.

Three problems:

1. infrared divergences.
2. ultraviolet divergences.
3. divergent series (even if each integral converges, the series of them may not).

The obvious thing to do is give up, which is what people historically did at first.

Why does the integral $\int_{-\infty}^{\infty} f(x)dx$ not converge? There are two basic reasons. (1) f may be locally bad (e.g. $f(x) = x^{-10}$, which blows up at zero), which is called an *ultraviolet divergence* (it is short distance singularity). The other reason an integral might not converge is (2) f is globally too large (e.g. $f(x) = x^2$), which is called an *infrared divergence*. In general, these integrals will have both kinds of divergences.

(1) Dealing with infrared divergences. The solution is to ignore them. The key point is that individual Feynman diagrams have infrared divergence, but if you sum over all Feynman diagrams of a given order, the infrared divergences automatically cancel out. There is a simple physical reason you'd expect this to happen. Suppose we're looking at the Lagrangian $m^2\phi^2 + (\partial\phi)^2 + \lambda(x)\phi^4$ for some function λ on spacetime.

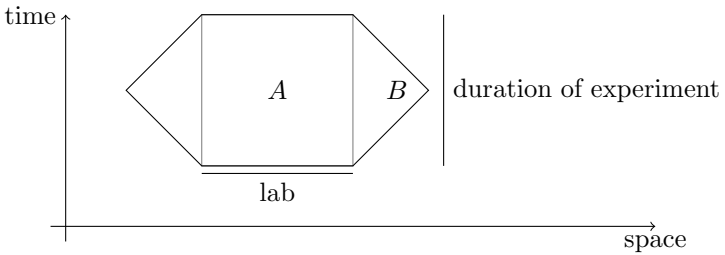

A=lab during experiment. It doesn't matter what happens what happens outside of the box A. Well it does matter because something can leave the lab and then come back. Let B=all points where you can send a signal and get it back. We don't care what goes on outside the region B. So there is some compact region which is all we care about. If we want, we can just replace the coupling coefficient λ by a function vanishing outside of B. If we do that, then we find that all the integrals are over a compact set, so there are no large-distance singularities.

Q: what do you mean by "ignoring them"? Do you pretend they are zero, or keep them around until they cancel? RB: The infrared divergence will be given by some parameter going to zero. The integral will diverge if you let this parameter go to zero. However, if you add all the integrals and then let the parameter go to zero, you get a finite number. Q: does it all go away, or do you get a non-zero finite term? RB: there is a non-zero finite term.

(2) Ultraviolet divergences. These are much trickier to deal with and they definitely don't cancel out. Distributions t have no ultraviolet divergences (almost by definition, $\int t(x)f(x)dx$ is defined for all f smooth of compact support). So we could get rid of ultraviolet divergences if we could replace $\Delta(x_1, x_2)\Delta(x_2, x_3)^3\Delta(x_3, x_4)$ by a distribution, we could eliminate ultraviolet divergences. Each of the factors is a distribution, but the product of distributions need not be a distribution (e.g. if you square the Dirac delta, you get nonsense).

Regularize the propagator $\Delta(x_1, x_2)$. This means you add an extra variable $\Delta(x_1, x_2, \varepsilon)$. For example, $\Delta(x_1, x_2)$ might be a Fourier transform of $(p^2 + m^2)^{-1}$. If you integrate over large values of p, it will be infinite. What you can do is replace $(p^2 + m^2)^{-1}$ by $(p^2 + m^2)^{-1-\varepsilon}$. For ε large and positive, everything converges nicely and we can define all Feynman integrals. Then we want to take $\varepsilon = 0$. What we do is we get an analytic function of ε and take the analytic continuation to $\varepsilon = 0$ and look at the value there. It turns out you can't do this because there is a pole at $\varepsilon = 0$ in general. You deal with this pole by renormalization.

3 RB 09-11

Today we'll try to finish off the first lecture. We said last time the infrared divergences all cancel out. You deal with ultraviolet divergences by regularization and renormalization. Regularization is where you raplace the propagator (say $\frac{1}{p^2+q^2}$) by, say, $\frac{1}{(p^2+q^2)^{1+\varepsilon}}$. Everything is a function of ε with poles at ∞. Renormalization is a cunning way to choose Lagrangian as a function of ε.

The third problem we had was divergent series. The problem is that we only know how to define Feynman integrals of the form $\int\int(\text{poly in }\phi)e^{\text{quadratic}}D\phi$, but we want to be able to define things like $e^{\text{non-quadratic}}$. The non-quadratic Lagrangian will be something like $(\partial\phi)^2+m^2\phi^2+\lambda\phi^4$. You can expand this as a formal power series in λ to get $\int e^{\text{quadratic}}D\phi+\int\int\lambda\phi^4e^{\text{quadratic}}D\phi+\int\int\lambda\phi^4\int\lambda\phi^4e^{\text{quadratic}}D\phi+\cdots$. Each of these integrals can be defined useing regularization and renormalization. So you get a well-defined formal power series in λ, which (probably) doesn't converge (unless $\lambda=0$). Nobody knows for certain that it doesn't converge in four dimensions, but it is ridiculously improbably. It doesn't converge because the coefficient of λ^n is the sum on Feynman diagrams on n points. The number of Feynman diagrams on n points grows like $n!$ or something. So we get series like $\sum\lambda^n n!$ (this shouldn't really be taken seriously), which doesn't converge for any non-zero λ. For λ small, the first few terms decrease, so you get a very accurate approximation to something (but nobody knows what). In QED, λ is something like 10^{-2}, so the first hundred terms decrease and then the terms are incredibly tiny, so you have a very good estimate of something, and these agree with experiment to great accuracy. This final problem nobody knows what to do about.

There is no caninical way to regularize and renormalize Feynman integrals. I vauguely indicated that you can insert ε in a particular way, but there are loads of other way to do it. There are something like a dozed different ways it's done in the litiriture. This corresponds to the fact that Feynman measure is not unique.

Suppose we have two different regularization/renormalization methods A and B, and suppose we have a Lagrangian L. This means we get two different theories for the Lagrangian L, which don't give the same answer. In other words, the QFT we get depends on L AND on the choice of renormalization method. The space of Lagrangians is well-behaved (it's finite-dimensional), but the space of renormalization schemes is not. If your theory depends on an unknown point in an infinite-dimensional space, it's really bad beacause there is no way to identify this point with a finite number of measurements.

Fortunately, there is a group of (finite) renormalizations which acts simply transitively on the space of renormalizations (i.e. Feynman measures), and it also acts on the space of Lagrangians. The theory of (L,M) (Feynman measure M) is the same as the theory of $(g(L),g(M))$ where g is a renormalization. So physicist A may be using a Lagrangian L, but phisicist B has to use a slightly different Lagrangian to get the same results. Suppose A uses (L,A) and B uses measure B. Find a renormalization g taking A to B and then the second physicist has to use Lagrangian $g(L)$. So there is no such thing as a theory associated to a Lagrangian, you need a Lagrangian together with a Feynman measure.

Warning 3.1. The group of renormalizations is not the same as the renormalization group, though they are closely related. The first is infinite-dimensional and acts on the space of Lagrangians and measures. The second is one-dimensional and is the group of "rescalings" of space-time. Rather confusingly, the second group also acts on measures and Lagrangians. Suppose you choose a measure A. By rescaling space-time, you get a new measure A'. You might think that A' is obtained by scaling A by a constant, but renormalization and regularization are not scale-invariant, so the relationship is more complicated than you'd think. These two measures are related by a renormalization. In other words, if you take a measure A and an element of the renormalization group, this gives you a renormalization. This isn't actually a homomorphism, but the way. ⌟

That's it for introduction. Let's get to producing examples and seeing what axioms they satisfy.

First recall representation theory of a Heisenberg algebra. It turns out that the theory of free quantum field theory is equivalent to the representation theory of Heisenberg algebras. The Heisenberg algebra is a Lie algebra of dimension 3 with $[X,Y]=Z$ and Z in the center. A typical example is $X=\frac{d}{dx}$ and $Y=x$ and $Z=1$ acting on $\mathbb{C}[x]$. We want to find all "lowest weight" representations of the Heisenberg

algebra generated by "vacuum vector" v with $Y(v) = 0$ and $Z(v) = v$. If you've done quantum mechanics, this is the harmonic oscillator. Y is the annihilation operator and X is the creation operator. It is easy to work out what the representation looks like. The only interesting thing you can do is act on v by X.

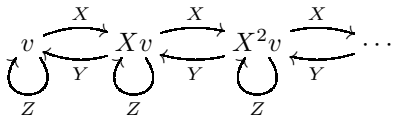

So the representation we get is the symmetric algebra $H = \mathbb{C}[x]$.

We want to put an inner product on H so that $X^* = -Y$ and $(v, v) = 1$. From this you can work out everything else. For example, $(Xv, Xv) = (v, -YXv) = (v, v) = 1$. It turns out that

$$(X^n v, X^n v) = n!(v, v).$$

PT: do you care if it's positive definite? RB: No, I don't care. The point is that the representation theory of Heisenberg algebras is easy (at least if you only care about highest weight representations). The same thing happens if instead of X and Y you take an arbitrary vector space.

Let V be a vector space with some inner product $(,)$. For a Heisenberg algebra $V^- \oplus \mathbb{R} \oplus V^+$, where $V^\pm$ are copies of V, with $[V^+, V^+] = 0 = [V^-, V^-]$ and $[u^+, v^-] = (u, v)$. Highest weight rep: acting on $S(V^+)$ with $V^-(1) = 0$ and $\mathbb{R}$ acts as multiplication by the identity. So the free quantum field theory we're going to construct is essentially a highest weight rep of a Heisenberg algebra.

Basic example: Free Hermitian scalar field. To construct a QFT, we need to give (1) a $*$-algebra A and (2) a $*$-representation H. A is generated by operators $\phi(f)$ where f is a classical field with compact support. "classical field" just means a smooth real function on space-time. $\phi(f) = \phi^+(f) + \phi^-(f)$ for certain operators $\phi^+(f)$ and $\phi^-(f)$ which are part of a Heisenberg algebra. The commutation relations are $[\phi^+(f), \phi^+(g)] = 0 = [\phi^-(f), \phi^-(g)]$ and $[\phi^+(f), \phi^-(g)] = i \int_{\mathbb{R}^n} \tilde{f}(p)\tilde{g}(-p)m(p)d^n p$, where the $\tilde{f}$ and $\tilde{g}$ are Fourier transforms of f and g, and $m(p)$ is the measure with support $p^2 = m^2$

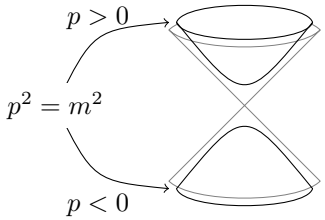

$m(p)$ supported on top sheet and is rotationally invariant. Take H to be a lowest weight representation of a Heisenberg algebra. there is a (very degenerate) inner product, whose kernel we quotient out by to get H. Then A is the algebra of operators generated by $\phi(f) = \phi^+(f) + \phi^-(f)$. $\phi(f)^* = \phi(f)$. Now we have all the basic data.

PT: why don't you use the whole Heisenberg algebra, just these combinations? RB: because the Wightman axioms don't care about them; you can keep them if you like.

Any unitary representation of a group is going to satisfy these conditions, so we need to narrow our definition of a QFT. What extra conditions do A and H need to satisfy to be a QFT? Answer: Wightman axioms.

1. (minor) The algebra A is generated by $\phi(f)$ where f is a classical field on space-time with compact support.

2. The inner product on H is positive definite. If you're an analyst you'll probably want to complete to get a Hilbert space, but there is no need for that.

3. Lorentz invariance: $O^+_{1,3}(\mathbb{R})$ (the + means preserving the time direction)[1] acts on $\mathbb{R}^{1,3}$, and A, H should be invariant under this action. In fact, this is Poincaré invariant because you also require translation invariance.

4. (Positive Energy condition) E =translation (forward) in time is a positive operator (i.e. $(Ea, a) \geq 0$ for any $a \in H$).

[1] Switching the time direction would switch which sheet of the hyperboloid you're on.

5. (Locality) $[\phi(f), \phi(g)] = 0$ if the supports of f and g are spacelike separated (i.e. if $f(x) \neq 0$ and $g(x) \neq 0$, then $x - y$ is spacelike). This is non-trivial to check in our example.

6. (Vacuum vector) There is a vector fixed by the Lorentz group $\mathbb{R}^{1,3}O^+_{1,3}(\mathbb{R})$.

7. (minor) the vacuum is essentially unique

8. (minor) Tempered.

4 RB 09-18

In the last lecture, we were discussing the free Hermitian scalar quantum field theory. Recall that, as for all quantum field theories, we need to give a $*$-algebra A acting on H. For a free quantum field theory, A is generated by operators $\phi(f)$ where f is a smooth compactly supported function on $\mathbb{R}^{1,3}$. You define $\phi(f) = \phi^+(f) + \phi^-(f)$, creation and annihilation operators. All the creation operators commute, and all annihilation operators commute, and $[\phi^-(f), \phi^+(g)]$ is some scalar ($\int_{p^2=m^2, p_0>0} \hat{f}(p)\hat{g}(-p)$). And $\phi^-(f)^* = \phi^+(f)$. This more or less gives you a Heisenberg algebra, and these algebras have really easy representation theory. H is the (essentially unique) representation generated by the vacuum vector which is killed by all $\phi^-(f)$.

I was commenting on various properties that this has: positivity, Lorentz invariance, positive energy. These are easy to check. A harder property to check is locality: $[\phi(f), \phi(g)] = 0$ if f and g are spacelike separated (i.e. Supp(f) and Supp(g) are spacelike separated). This has something to do with the fact that you can't send signals faster than light.

We have

$$\begin{aligned}[\phi(f), \phi(g)] &= [\phi^-(f), \phi^+(g)] - [\phi^-(g), \phi^+(f)] \\ &= \int_{p^2=m^2, p_0>0} \hat{f}(p)\hat{g}(-p) - \int_{p^2=m^2, p_0<0} \hat{f}(p)\hat{g}(-p) \\ &= \int f(x)g(y)m(x-y)\, d^4x\, d^4y\end{aligned}$$

[[★★★ picture of cone and two sheet hyperboloid; first integral over top sheet and second integral over bottom sheet.]] This is the same as integrating over the whole two sheet hyperboloid with a measure M of coefficient $+1$ on the top sheet and a coefficient of -1 on the bottom sheet. Let m be the Fourier transform of this measure. So to prove locality, we have to show that m vanishes on spacelike vectors. You could explicitly compute m, but let's not. M has the following properties:

1. It is invariant under rotations preserving the time direction.

2. It changes sign under the reflection $t \mapsto -t$.

It follows that m has the same symmetries. Any function m with these symmetries is zero on spacelike vectors x because there is a time-reversing rotation σ fixing x. So $m(x) = m(\sigma x) = -m(x)$ as m changes sign under σ. With a little more work, we see that any measure m must also vanish on spacelike vectors.

So locality can be traced back to the fact that the measure we integrate over has an antisymmetry property under reversing the time direction.

Why do we define $[\phi^-(f), \phi^+(g)] = \int_{m^2=p^2, p_0>0} \hat{f}(p)\hat{g}(-p)d\cdots$? (1) We want translation invariance. Any translation invariant distribution like this can be written as $\int \hat{f}(p)\hat{g}(-p)d(\text{measure})$ by taking Fourier transforms and fiddling with them. Conversely, you can take any measure you like and this will be translation invariance. (2) We want rotation invariance. This forces the measure to be rotationally invariant. (3) Positive energy condition needs the condition that the support of the measure has "positive energy". This means that the support is in the positive cone of momentum space ($p^2 > 0$). [[★★★ only the top part of the cone for some reason]] (4) Why do we need the condition $p^2 = m^2$? We don't. For other measures, we get things called "generalized free field theories" which are pretty similar. The only use I know of for these things is as a source of counterexamples. Q: the ones for fixed m are irreducible and the others aren't? RB: depends on what you mean by that. It's true that if you want the state space (?) to be irreducible, you have to use the free field theory. The generalized guys still have H irreducible as an A-representation.

What is a free field (theory)? That is, how do you recognize a free field theory when you see it? There doesn't seem to be any definition, but everybody can recognize one when it comes up.[1] A free field theory should be roughly equivalent to a representation of some Heisenberg algebra. The $*$-algebra should be generated by a set of "annihilation operators" which commute (this isn't too important) and the commutator of an annihilation operator with an adjoint of one should be some scalar (this is the key property).

What does a free field theory actually look like? That is, what is the Hilbert space H? $H = \mathbb{C} \oplus Sym^*(H_1)$ (H_1 is "one particle states"). H_1 is the space of (well-behaved)[2] functions on the manifold $p^2 = m^2, p_0 > 0$. Classically, the one particle states with $p^2 = m^2$ correspond to the points of this manifold. In quantum mechanics, the space of one particle states is functions on this manifold. The interpretation is that $Sym^n(H_1)$ is the space of n non-interacting particles ("sort of like mathematicians at a party I guess"). It would be good to have a field theory where you have the particles interacting. Nobody has found such a field theory which satisfies the Wightman axioms. The problem with the Wightman axioms is that they don't allow perturbative QFTs, where instead of working over $\mathbb{R}$, you work over a formal power series ring $\mathbb{R}[[\lambda]]$ (λ is called a coupling constant). Solution: extend the Wightman axioms to work over $\mathbb{R}[[\lambda]]$; make A an algebra over the power series ring instead of over $\mathbb{R}$. Let's check if there are any problems.

- Lorentz invariance works over $\mathbb{R}[[\lambda]]$ with no changes.
- Locality works over $\mathbb{R}[[\lambda]]$ with no changes (so long as you don't do anything stupid).
- Positive energy and positivity gives a problem. $(a, a) \geq 0$ and $(a, Ea) \geq 0$. You have to decide what it means for a formal power series to be positive. There are several different ways to define positivity. We need to explicitly choose a set of positive elements of $\mathbb{R}[[\lambda]]$. It doesn't seem to matter very much which method you choose. For example, you could say $x \geq 0$ if x is a square of a formal power series (any sum of squares is already a square). AJ: does this admit $\Phi 3$(?) theory in 6-dimensions, where you don't have a positive definite Hamiltonian by perturbatively it doesn't matter? RB: I'm not sure. The problem is that the potential is not positive, but perturbatively it is (because the kinetic term is positive). I would guess this works. This illustrates that working perturbatively, you don't see a lot of important things.

To describe a free field theory, we'll describe its n-point "functions" (distributions). To describe a representation H from A using a "state" ω, we do $(a, b) = \omega(ab^*)$. What is the state ω for a free field theory? A typical element of A looks like $\phi(f_1)\cdots\phi(f_n)$. The corresponding element ω is given by $\langle vac, \phi(f_1)\cdots\phi(f_n)vac\rangle$. This is sometimes written as $\langle vac|\phi(f_1)\cdots\phi(f_n)|vac\rangle$ or as $\langle\phi(f_1)\cdots\phi(f_n)\rangle$ [[★★★ I think]]. This is really a distribution on the product of n copies of spacetime (sometimes

[1] A dog can't define a rat, but knows one when it smells it.

[2] L^2 or rapidly decreasing, or whatever.

called a Wightman distribution). Knowing this ω is the same as knowing these Wightman distributions explicitly.

The idea is that to compute $\langle vac, \big(\phi^+(f_1) + \phi^-(f_1)\big) \cdots \big(\phi^+(f_n) + \phi^-(f_n)\big) vac\rangle$, you push the $\phi^-(f_i)$ to the right, as these kill the vacuum. Let's do the case of a two-point function first.

$$\begin{aligned}\langle(\phi^+(f_1) + \phi^-(f_1))(\phi^+(f_2) + \phi^-(f_2))\rangle &= \langle\phi^-(f_1)\phi^+(f_2)\rangle \\ &= \langle[\phi^-(f_1), \phi^+(f_2)]\rangle\end{aligned}$$

which is a scalar distribution given by $\Delta^+(x_1, x_2)$, Fourier transform of $p^2 = m^2$, $p_0 > 0$. This is some kind of Bessel function. [[★★★ rewrite this. since the commutator $[\phi^-(f_1), \phi^+(f_2)] =: \Delta^+(x_1, x_2)$ is a scalar operator, $\langle vac, [\phi^-(f_1), \phi^+(f_2)] vac\rangle = [\phi^-(f_1), \phi^+(f_2)]\langle vac, vac\rangle = [\phi^-(f_1), \phi^+(f_2)]]]$

Lemma 4.1. *The "n-point function" $\langle\phi(f_1) \cdots \phi(f_n)\rangle$ is given by $\sum_{\sigma \text{ perfect pairing}} \prod_{i<\sigma(i)} \Delta^+(x_i, x_{\sigma(i)})$, where a* perfect pairing *is an order 2 permutation with no fixed points.*

Proof. It suffices to prove that

$$\langle\phi(f_1) \cdots \phi(f_n)\rangle = \sum_{k=2}^{n} \Delta^+(x_1, x_i) \cdot \langle\widehat{\phi(f_1)}\phi(f_2) \cdots \widehat{\phi(f_k)} \cdots \phi(f_n)\rangle$$

(where the hats indicate omission) and to observe that the 1-point function $< vac, (\phi^+(f_1) + \phi^-(f_1)) vac\rangle$ is zero because the ϕ^- kills the vacuum on the right and the ϕ^+ kills the vacuum on the left. This formula follows almost immediately from the observation

$$\begin{aligned}\phi^-(f_1)\phi(f_k) &= [\phi^-(f_1), \phi^+(f_k)] + \phi(f_k)\phi^-(f_1) \\ &= \Delta^+(x_1, x_k) + \phi(f_k)\phi^-(f_1).\end{aligned}$$

[[★★★ it wouldn't hurt to write this out more explicitly]] □

Example 4.2. By the lemma, we can compute the 4-point function $\langle\phi(f_1)\phi(f_2)\phi(f_3)\phi(f_4)\rangle$ to be

$$\begin{aligned}&\Delta^+(\,1 \text{—} 2 \quad 3 \text{—} 4\,) + \Delta^+(\,1 \overset{\frown}{\quad 2 \quad 3 \quad} 4\,) + \Delta^+(\,1 \overset{\frown}{\quad 2 \text{—} 3 \quad} 4\,) \\ =&\Delta^+(x_1, x_2)\Delta^+(x_2, x_4) + \Delta^+(x_1, x_3)\Delta^+(x_2, x_4) + \Delta^+(x_1, x_4)\Delta^+(x_2, x_4)\end{aligned}$$ ◇

So in free field theories, everything can be determined in terms of the two-point functions. This two-point functions is called the Wightman 2-point function or the *cut propagator*.

Next we'll discuss propagators, which are confusing because there are six different sorts of propagators for a QFT. We'll talk about them and the relations between them.

5 RB 09-25

Propagators

If you look in a physics book, a propagator is supposed to be an "amplitude for a particle to go from one point to another". I don't know that this means, so I'll ignore it. Six types of propagators (for the hermitian scalar field theory)

1. 2 Feynman propagators
2. 2 Cut propagators
3. 2 advanced/retarded propagators

If you think that's a lot, each propagator can be viewed in position space, but it also has a Fourier transform which lives in energy/momentum space, so all together, there are twelve things. Moreover, propagators can be either *massless* or *massive*, so there are 24 propagators. You also have to worry about spin 0, $\frac{1}{2}$, or 1 propagators, so there are 72 of them. These are propagators in Lorentzian space. We could also look a propagators in Euclidean space (his doesn't quite double it). We can also look at propagators in other dimensions, so you should multiply this 72 by ∞. In dimension ≥ 2, massless propagators have special properties. They also behave differently in odd or even dimensions.

That gives you a bird's-eye view of propagators. What is a propagator Δ? Let's consider the case of a Hermitian scalar field in 4-dimensional Lorentzian space.

(1) Δ is a distribution on $M \times M$, where M is spacetime. If you think of Δ as a function (which it isn't), $\Delta(x, y)$ is the "amplitude of propagation from x to y".

(2) It is translation invariant: $\Delta(x + z, y + z) = \Delta(x, y)$, so it is much easier to think of it as a distribution on one variable: $\Delta(x, y) = \Delta(x - y)$ is now a distribution on M.

(3) Δ is a solution to $(\partial_i^2 + m^2)\Delta(x) = c\delta(x)$ (this first operator is the Klein-Gordan operator), where $\delta(x)$ is just a Dirac delta function at 0. If $c = 1$ (Feynman, advanced/retarded), Δ is a Greens function for the K-G equation. If $c = 0$ (Cut propagators), Δ is a solution of the K-G equation.

(4) Δ is invariant under rotations (preserving time).

(5) Wave front set of Δ should be as small as possible (see later).

These are the properties that characterize the propagators we'll be interested in. Before we find some propagators, let's say what they are used for.

1. Cut propagators appear as the two point functions of a free field theory. More generally, all the n-point functions can be written in terms of the cut propagator $\Delta(x, y)$.
2. They are needed to define Feynman diagrams. A Feynman diagram is going to appear as a piece of an asymptotic expansion.

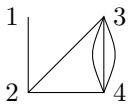

Each edge in the diagram represents a Feynman propagator, and the diagram represents their product, so the picture above represents the term $\Delta(x_1, x_2)\Delta(x_2, x_3)\Delta(x_2, x_4)\Delta(x_3, x_4)^3$. This is NOT DEFINED because of ultraviolet singularities.

Let's try to solve $(\partial_i^2 + m^2)\Delta = \delta(x)$ (i.e. Δ is a Greens function for the K-G operator). Why are we interested in the K-G operator, by the way? Because it appears as the E-L equation for the "free" part of the Lagrangian. For this lecture, assume somebody has given you the K-G operator.

First let's solve it in Euclidean space where it is rather easier. Let's try to solve $(\partial_i^2 - m^2)\Delta = \delta(x)$ with Δ tempered (i.e. don't behave too badly, so closed under Fourier transforms). Taking the Fourier transform, we get $(p^2 + m^2)\tilde{\Delta} = 1$ (where $p^2 = p_1^2 + p_2^2 + \cdots$). The solution is that $\tilde{\Delta} = 1/(p^2 + m^2)$. Note that $p^2 + m^2 > 0$, so this is well-defined. So Δ is the Fourier transform of this $\tilde{\Delta}$. You can write this as a Bessel function if you want, but this doesn't completely specify it as a distribution because of something with singularities. So this is the only reasonable thing to use as a propagator in Euclidean space, so we don't get 6 different things. Note that if $m = 0$ you have some problems, but let's not worry about it.

Now let's try it in Lorentz space. As before, we get $(p^2+m^2)\tilde{\Delta}=1$. Now we have a problem because $p^2+m^2=0$ on a two sheet hyperboloid (or a cone if $m=0$). Anywhere off these surfaces, you can invert p^2+m^2. How do we define $1/(p^2+m^2)$ (as a distribution) when $p^2+m^2=0$? First of all, the solution of $(p^2+m^2)\tilde{\Delta}=1$ is not unique. This equation has solutions given by any function (times measure) on the hypersurface $p^2=m^2$ (more precisely, any distribution on this hypersurface gives you a solution). So there is an infinite-dimensional space of solutions. We can cut down the dimension by adding the condition that Δ is invariant under time-preserving rotations. In the Euclidean case, it was automatically invariant under rotations, but in the Lorentzian case, this really is a new condition. The rotation invariant solutions of $(p^2+m^2)\tilde{\Delta}=0$ form a 2-dimensional space (one for each component of $p^2+m^2=0$). No, that's a bit misleading. This 2 is the number of orbits of the group of time-preserving rotations on the space $p^2=m^2$; this is not the number of topological components.

Remark 5.1. If the dimension is 2 and $m=0$, then $p^2=0$ has FIVE components under the connected component of the rotation group.

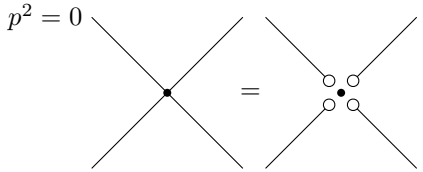

This corresponds to left moving and right moving things on a string, so this shows up in string theory. ◇

The cut propagators are the two solutions to $(p^2+m^2)\tilde{\Delta}=0$ corresponding to these two components. So a cut propagator is the Fourier transform of the invariant measure on a component of $p^2=m^2$.

Now we want some "good" solutions to $(p^2+m^2)\tilde{\Delta}=1$. The solution is essentially $1/(p^2+m^2)$, but we have to explain what this means when $p^2=-m^2$ [[There is a sign convension which I can't ever rememeber]]. There are four reasonable things it could mean. The Fourier transform of $1/(p^2+m^2)$ is $\int_{\mathbb{R}^4}\frac{e^{ipx}}{p^2+m^2}d^4p$. when $p^2=m^2$, this blows up, but if you think of p as complex, then there is a pole at $p^2+m^2=0$, and you can go around the pole.

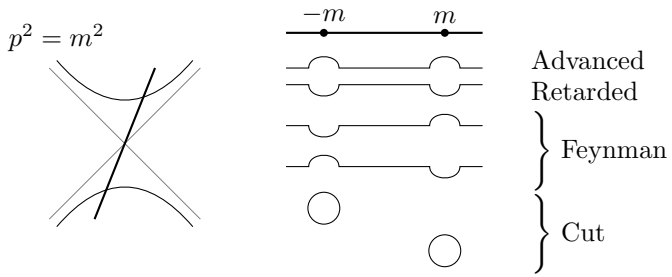

There are four reasonable ways to go around the poles. These four choices give the remaining propagators. To get the cut propagators, you're really taking the residue at the pole (i.e. integrating around the little circles). Notice that you can read off lots of linear relations from linear relations between these 1-cycles. For example, if you take the difference between the advanced propagator and a Feynman propagator, you get a cut propagator. The space of solutions is only three dimensional, so if you take any four of these, there will be a linear relation among them.

We're integrating e^{ipx}. This integral goes to zero if the imaginary part of p goes to infinity. So if you have an advanced propagator, you can move your line of integration up to get that it vanishes. The conclusion: The advanced and retarded propagators vanish if x is not in some closed cone.

The advanced and retarded propagators aren't used very often, except to show that some Feynman propagator is equal to some cut propagator except inside a certain cone.

What is special about these six propagators? They have very small WAVE FRONT sets. I said that there is a 3-dimensional space of Greens functions. Why not integrate along some path that loops around the poles a few times? That gives you a perfectly good propagator, but they have large wave front sets.

Motivation: we want to multiply propagators together (because we want to make sense of Feynman diagrams). The problem is that we

cannot multiply distributions. For example, there is no reasonable way to make sense of $\delta(x)\cdot\delta(x)$. We can multiply distributions with DISJOINT SINGULAR SUPPORT because you can always multiply a distribution by a smooth function. The problem is that for Feynman diagrams we want to multiply propagators without disjoint singular support. Wave front sets give a more refined obstruction to multiplying distributions.

Example 5.2. $\delta(x)\cdot\delta(x)$ =?. However, (think of δ as eating functions and giving values) $\delta(f) = f(0)$ which is the residue of $f(z)/z$ at zero, which is given by the integral $\int_C \frac{f(z)}{z}dz$, where C goes around 0 counterclockwise. Alternatively, we could integrate $\frac{1}{2\pi i}\int_{-\infty}^{\infty}\frac{f(z)}{z}dz$ in two different ways (going a little above and a little below the origin), and the difference of these two give you $\delta(f)$. So $\delta(x) = \frac{1}{2\pi i}\frac{1}{x^+} - \frac{1}{2\pi i}\frac{1}{x^-}$. These distributions $\frac{1}{x^\pm}$ are equal to $\frac{1}{x} = \frac{d}{dx}\log(x)$ for $x \neq 0$. $\frac{1}{x^+}$ is $\frac{d}{dx}\log(x)$ for the imaginary part of x greater than zero.

Now note that $\left(\frac{1}{x^+}\right)^2$ is perfectly well defined because it is the boundary value of $\frac{1}{x}$ on the upper half plane. However, $\frac{1}{x^+}\times\frac{1}{x^-}$ is not well-defined because you cannot multiply functions on the upper and lower half planes. The point is that the wave fronts of these two are small enough. ◇

Next time we'll say what wave fronts are and why you can use them to tell if you can multiply propagators.

6 RB 10-02

Last lecture we looked at the following problem: (1) a Feynman diagram represents a product of propagators (one for each line). The problem is that propagators have singularities, which make it difficult ot multiply them. There is one easy case when you can multiply them (when they have disjoint singular supports), but this is not refined enough for us. We saw that you can sometimes multiply two distributions even when they are singular in the same place.

Question: when can you multiply two distributions f and g at a point x? Answer: whenever the wave front sets at x travel in roughly the same directions.

What is a waved front set of a distribution f? A crude measure of singularities: look at singular points of f. The wave front set is a subset of the cotangent space at each singular point. PT: how are you thinking of these distributions? As functions which are allowed to blow up? RB: yes, that is a good approximation, and the singular set is the set of points where f cannot be written as a smooth function. A compactly supported distribution f on $\mathbb{R}^n$ is smooth (at all points) is equivalent to saying that the Fourier transform $\hat{f}$ is rapidly decreasing. Now we can ask, "in which directions is $\hat{f}$ NOT rapidly decreasing?" These directions somehow tell you the "directions in which the singularities are going". Such a direction is an element of the dual of $\mathbb{R}^n$.

Definition 6.1. The *wave front* set of f at $x \in \mathbb{R}^n$:

1. localize f at x: multiply f by a smooth compactly supported bump function u with $u(y) = 1$ for y near x.
2. look at the Fourier transform $\widehat{fu}$ of fu. This is a function on the cotangent space of x

The wave front set of f at x is given by intersection over all u of the directions in T^*_x near which $\widehat{fu}$ is NOT rapidly decreasing. ◇

This makes the wave front set of f a conical subset of the cotangent space of $\mathbb{R}^n$

The singular points are just those such that there is a non-zero element of the wave front set in the cotangent space at that point. Some authors

say you shouldn't count the zero covector as being in the wave front set, but sometimes it is handy to include the zero covectors.

We can multiply f and g provided there is no point x and vectors $v \in WF(f)_x$ and $w \in WF(g)_x$ with $v + w = 0$ with $v, w \neq 0$. $WF(f)_x$ is a cone and $WF(g)_x$ is a cone. If the cone generated by $WF(f)_x$ and $WF(g)_x$ is a proper cone, then we can multiply f and g. This also works for collections of distributions [[★★★ because the wave front set of the product is contained in the cone generated by the wave front sets of the factors?]].

What are the wave front sets of our propagators? There are six different propagators because there are six interesting choices of wave front sets. The singular points all lie on norm zero vectors (points in the usual cone). We can think of the cotangent space as the tangent space, which inherits the metric. It turns out that the vectors of the wave front sets are cotangent vectors of norm zero (they lie on the cone *in the cotangent space*).

Possible wave front vectors p: $(x > 0 \text{ or } x < 0) \times (p > 0 \text{ or } p < 0)$ (there is also $x = 0$, but that is a really bad guy). We can tweak our propagators so that that the wave front set vanishes in two of these four regions. This gives us $\binom{4}{2} = 6$ possibilities.

(1) Advanced, Retarded propagators. These had the property that the support is contained in a closed cone, so all the singularities are in that cone, so these are the cases where the wave front sets lie in $(x < 0)$ or $(x > 0)$. These don't occur in quantum mechanics because the wave front sets where these are singular are going in all directions, so you can't multiply them together on the light cone. So advanced and retarded propagators are good for classical mechanics, but not for quantum.

(2) Cut propagators. These have the property that the Fourier transforms $\widetilde{\Delta}$ have support on one of the two hyperboloids of revolution of $p^2 = -m^2$ (i.e. one of the sheets of the usual hyperboloid). This is rapidly decreasing except in directions $p^2 = 0$ with $p > 0$ (or $p < 0$) (this is because we're supported on the sheet, so in any other direction, you're eventually zero). So the wave front sets have $p > 0$ (or $p < 0$). Note that you can multiply a cut propagator by itself as much as you want because the wave front sets always go in the same direction.

(3) Feynman propagators Δ. In this case, Δ is equal to a cut propagator except in the negative cone (since Δ is a cut propagator plus the retraded propagator), so Δ has the same singularities as the cut propagator except in the negative cone. However, Δ is equal to the other cut propagator except in the positive cone (using another relation), so in the positive cone it has the same singularities as the other cut propagator.

Possible wave front sets of propagators:

[[★★★ picture]]

We can multiply Feynman propagators by themselves everywhere except at zero, which is why we get ultraviolet divergences.

Now let's try to evaluate some Feynman diagrams: put a Feynman propagator at each edge and try to multiply them together. By the way, we haven't specified the directions of the edges, but since the Feynman propagators are invariant under multiplication by -1, we're fine.

[[★★★ x y connected by two edges]]

This is $\Delta(x - y)^2$, which is defined except at $x = y$. If we require translation invariance, this implies that the ambiguity is a distribution on $\mathbb{R}^n \times \mathbb{R}^n/\mathbb{R}^n$ with support at a point, given by the diagonal mod $\mathbb{R}^n$. Distributions supported at a point are really easy to deal with.

[[★★★ complicated picture]]

Suppose not all points are the same, and suppose it is connected. We can choose x and y which are joined by a line L. Let F be the diagram minus this line, so the diagram is $F \cup L$. A distribution of $F \cup L$ should be the distribution of F times $\Delta(x - y)$. When is this well defined? Check it is well defined by looking at the wave front sets. For this we need to know about the wave front set of a diagram, which we should be able to do inductively.

Suppose $(p_1, \ldots, p_k)$ is in the wave front set of a Feynman diagram at the point $(x_1, \ldots, x_k) \in (\mathbb{R}^n)^k$.

Theorem 6.2. *If x_i is a minimal*[1] *point with $\sum_{x_j = x_i} p_j \neq 0$, then $\sum_{x_j = x_i} p_j < 0$. If x_i is a maximal point with $\sum_{x_j = x_i} p_j \neq 0$, then $\sum_{x_j = x_i} p_j > 0$*

Wave front set of Δ over a typical point $(x_1, \ldots, x_n)$ (with $x_1 > x_2$) looks like $(p, -p, 0, 0, \ldots, 0)$, $p > 0$. Wave front set of a Feynman diagram F must contain $(-p, p, 0, \ldots, 0)$ for the product not to be defined, but

[1] Space time is partially ordered: $x < y$ means we can send a signal from x to y.

this contradicts the condition satisfied by the wave front set of a Feynman diagram (the theorem).

We also need to check that the wave front set of $F \cdot \Delta$ also satisfies this condition. For this you need to know the wave front set of a product. Use the fact that the wave front set of a product of distributions at a point is contained in the sum of their wave front sets. It is then a fairly easy exercise to check the condition [[★★★ HW: do this exercise, proving the Theorem]].

The result is that we can define each Feynman diagram *up to* addition of a distribution that is (1) supported on the diagonal and (2) translation invariant. This is effectively a distribution supported at a point (the point diagonal/translations in $(\mathbb{R}^n)^k$/translations). PT: so there are only problems when all the points of the Feynman diagram are the same? RB: That's right.

Distributions with support at $0 \in \mathbb{R}^n$ are just given by polynomials in $\frac{\partial}{\partial x_i}$ applied to the Dirac delta function. The Fourier transform is therefore a polynomial in $p_1, \ldots, p_n$. So the ambiguity in the result is a polynomial in momentum. It will turn out that the amibuity will be closely related to what physicists called "counterterms".

7 RB 10-09

Last lecture, we said what a Feynman diagram is. It represents some product of propagators. We saw that the product of propagators is well defined UP TO addition of a translation invariant distribution with support on the diagonal. The Fourier transform of a distribution with support on the diagonal will be a polynomial in momentum. In physics, you get these things called counterterms associated to Feynman diagrams, and the fact that they are polynomial essentially comes from the fact that the product of propagators is defined up to this distribution on the diagonal. One problem is that we have to specify the product precisely (this is the problem of renormalization). We won't do this today.

Why are we interested in defining Feynman diagrams? Answer: they come from expanding Gaussian integrals $\int \text{polynomial}(x) \cdot e^{-\text{quadratic}(x)} dx$. Whenever you have such an integral, you'll probably get Feynman diagrams popping up. Pretend you're a 1A student.

1-dimensional case. (1) what is $\int e^{-x^2} dx$ (we know it is $\sqrt{\pi}$, but let's say we don't; the tricks won't work in infinite dimensions). We're stuck. (2) What about $\int x e^{-x^2} dx$? This is easier because it is $-\int \frac{1}{2} \frac{d}{dx} e^{-x^2} dx = 0$ since e^{-x^2} vanishes at $\pm\infty$. (3) How about $\int x^2 e^{-x^2} dx$? Integrating by parts, we can reduce to the case $\int e^{-x^2} dx$.

In general, $\int x^n e^{-x^2} dx = -\frac{1}{2} \int (n-1) x^{n-2} e^{-x^2} dx$. Diagrammatically, we have

[[★★★ picture]]

The whole integra is the sum of $-\frac{1}{2}$ times what you get if you pair off two xs and cross them out (leaving $n-2$ other xs), and sum over all possible pairings. Consider the case of x^6, we get a sum over all Feynman diagrams like

[[★★★ picture]]

multiplied by the integral $\int e^{-x^2} dx$. We put the "propagator" $-\frac{1}{2}$ on each edge (this $-\frac{1}{2}$ is a distribution on 0-dimensional space).

Now let's try to make some Feynman diagrams on more points.

Example 7.1. $\int x^4 \times x^4 \times x^6 \times e^{-x^2} dx$. Pretend we're particularly stupid 1A students and we haven't realized we can multiply these together. Then we sum over the diagrams

[[★★★ picture]] ◇

In higher dimensional cases, you can't always stick all the points together.

Now let's try to work out $\int e^{-m^2\phi^2-\lambda\phi^4/4!}$, which is a kind of 0-dimensional version of $\int e^{\int (m^2\phi^2(x)+(\partial\phi)^2+\lambda\phi^4)dx}$ for $\phi\colon \mathbb{R}^n \to \mathbb{R}$. This is defined for $\lambda > 0$, and you can define it for all complex $\lambda \neq 0$ by changing $\phi \mapsto \phi \cdot \lambda^{-1/4}$, but it has a branch point AND an essential singularity at $\lambda = 0$. At the point $\lambda = 0$, the integral converges, but there is still an essential singularity (if you approach from positive λ, it's ok, but if you approach from negative λ, you run into trouble). Expanding as a power series in λ at $\lambda = 0$ is totally stupid because you can't expand an essential singularity like this, but we're going to do it anyway. The integral becomes $\sum \frac{(-\lambda)^n}{n!} \int (\phi^4/4!)^n e^{-m^2\phi^2} d\phi$. Each term is something we can expand using Feynman diagrams. It is

$$\sum_n \frac{(-1)^n}{(4!)^n n!} [[\bigstar\bigstar\bigstar \text{ picture}]]$$

with the propagator $1/2m^2$ for each edge, and a factor of λ for each vertex. There are a lot of such diagrams and we'd like to reduce the number a bit. Instead of summing over diagrams, we can sum over isomorphism classes of diagrams

$$\sum_{isoclasses} (-1)^n [[\bigstar\bigstar\bigstar \text{ squiggle}]] \frac{\#\text{ diagrams in isoclass}}{(4!)^n n!}$$

and this last factor is one over the size of the automorphism group of the diagram. The reason is that the denominator is the number of automorphisms of the vertices, and if you go to a corner and think about it you'll see that it's right. So we get that what we want is

$$\sum_{isoclasses} (-1)^n \times \frac{\text{value of diagram}}{\text{order of automorphism group}}$$

This kind of weighting a diagram by one over the size of its automorphism group occurs all over the place in mathematics.

Back to our example of $\int e^{-m^2\phi^2-\lambda\phi^4} d\phi$" = " $\int e^{-m^2\phi^2} d\phi \times$ [[$\bigstar\bigstar\bigstar$ picture]]

But the power series can't possibly converge because there is an essential singularity, so what is the meaning of this power series. The power series is an *asymptotic expansion* of the integral valid for $\lambda > 0$. This means that the integral is asymptotic to $a_0 + a_1\lambda + a_2\lambda^2 + \cdots$, so it is equal to $a_0 + \cdots + a_n\lambda^n + O(\lambda^{n+1})$. So the approximation (for a given n) gets better as λ goes to zero [[$\bigstar\bigstar\bigstar$ or infinity?]], but the approximation gets worse as you increase n.

Finite dimensional case is pretty similar: $\int P(x)e^{-Q(x)}dx$ where $x \in \mathbb{R}^n$, P is polynomial and Q is quadratic. The calculation is similar; you just need to compute the propagator for something like $\int x_i x_j e^{-Q(x)}dx$. It is the bilinear form Q^{-1} applied to x_i, x_j.

The infinite dimensional case is somewhat trickier. What is (say) $\int P(\phi)e^{-Q(\phi)}D\phi$, where ϕ is a field on $\mathbb{R}^3$. For example, you could have $\int \phi(x_1)\phi(x_2)\phi(x_3)^3 e^{-\int \phi(x)^2+(\partial\phi)^2 dx} D\phi$ (ignoring the complication that ϕ is really a distribution). Trying to expand like before, we run into a severe problem right away.

Problem 1: What is $D\phi$? it should be a translation invariant measure on an infinite dimensional space and this is a big problem because translation invariant measures tend to exist only on locally compact spaces and infinite dimensional spaces are not locally compact. Such a measure doesn't exist. It turns out it is possible to make sense of the Gaussian measure $e^{-Q(\phi)}D\phi$. There are two approaches.

(Analytic approach) On any (real) finite dimensional Hilbert space, we have a canonical Gaussian measure $e^{-\pi x^2}d^n x$ of total mass 1. These are all compatible in the following sense. If we have a finite dimensional vector space $H_1 = H_2 \oplus H_3$ and we have projection $H_1 \to H_1/H_2 = H_3$, then the projection of the Gaussian measure of H_1 to H_3 is the Gaussian measure of H_3. This is obvious from the simple example: $\int f(x)e^{-\pi(x^2+y^2)}\,dx\,dy = \int f(x)e^{-\pi x^2}\,dx$.

Now suppose that H is infinite dimensional. Looking at all finite dimensional quotients, we see that all the Gaussian measure on the quotients are compatible, so we can define the measure of any cylindrical set (the inverse image of a measurable set of some finite dimensional quotient). This ought to give us a nice Gaussian measure on all of H. If it did, quantum field theory would be easy. It is not obvious what goes wrong with this construction. (1) $\mu(H)$ should be 1. (2) Suppose B_r is a ball

of radius r in H, we get that $\mu(B_r)$ is less than or equal to the measure of a ball of radius r in $\mathbb{R}^n$, which is less than or equal to the measure of a cube of side $2r$ in $\mathbb{R}^n$, which is less than or equal to $\left(\int_{-r}^{r} e^{-\pi x^2}\, dx\right)^n$, which tends to zero. So a ball of radius r has measure zero. Since H is a countable union of such balls, we get that $\mu(H) = 0$.

So what is wrong? The problem is that we have a "measure" defined on all cylindrical sets, which don't form a σ-algebra. It turns out that you cannot extend this to the σ-algebra generated by cylindrical sets in a countably additive way.

You can define a Gaussian measure on larger spaces of distributions. You have to replace your Hilbert space by something called a rigged Hilbert space, and then do some other stuff. It turns out that this bigger space consists of distributions rather than functions and this is bad because ϕ^4 doesn't make sense for a distribution ϕ.

Next week well see an algebraic approach.

8 RB 10-16

Guassian measure on infinite-dimensional spaces

Finite-dimensional Hilbert spaces have nice Guassian measures. For some subtle technical reason, you can't just take a limit of the Guassian measure on the finite-dimensional subspaces to get a Guassian measure on an infinite-dimensional space (you just get a cylindrical measure which cannot be extended).

The problem is that cylindrical set measures on Hilbert space do not in general give you honest measures. There is a way around this. If you have a Hilbert space of functions, it will often come equipped with a *nuclear space* as a subspace and it will be contained in the dual of a nuclear space. Such a thing is called a *rigged Hilbert space*. A typical nuclear space is something like the space of smooth rapidly decreasing functions; a typical example of a Hilbert space is L^2 functions; and the dual of the nuclear space is just distibutions. We DO get gaussian measures from cylindrical measures on rigged Hilbert space.

In quantum field theory, you run into a new problem. We want to integrate things like $e^{-\int \phi(x)^4 dx}$. If ϕ is L^2, then you have some chance of integrating $\phi(x)^4$, but if ϕ is a distribution, then ϕ^4 will not make any sense at all (if you try to define it with limits, it becomes infinite almost everywhere). In LOW dimensions, this works.

In the 1-dimensional case, the support of the measure is the set of "brownian motion paths" (non-differentiable, but at least continuous, so ϕ^4 makes sense). The result of this is that quantum field theory is easy in dimension 1. This measure is called Wiener measure.

In the 2-dimensional case, it fails, jut "only just". Roughly, ϕ has logarithmic singularities everywhere (which are "only just" singular).

The nastyness of the measure depends on how bad the singularities of a propagator are. In one dimension, the Green function is $\sum(\frac{\partial}{\partial x})^2)$, giving you $|x|$ which is continuous. In two dimensions, you get $\log|x|$, which is just barely singular. In three dimensions, you get $|x|^{-1}$, which is not borderline at all.

If you work at it, you can make two-dimensional QFT work out. In three dimensions an higher, you can't get around it (actually, some specialized case has been worked out).

An honest attempt to define gaussian measure did not work, so we're going to cheat. If you can't solve a problem, secretly change the definitions in the problem to make it easier. We'll change the definition of a measure. A measure (1) assigns a real to every measurable set such that .... (2) Alternatively, we can use the idea of a Radon measure, which (for locally compact spaces) can be thought of as a linear map (compactly supported continuous functions)$\to \mathbb{R}$, thought of as $f \mapsto \int f d\mu$ such that .... These two definitions are more similar than they appear; you can think of a normal measure as a linear function from measurable functions to $\mathbb{R}$.

Generalization: define a measure to be a (well-behaved) map from some space of functions on X to $\mathbb{R}$. So we'll only worry about integrating some smaller set of functions which we're really interested in. A typical example of such a space of functions will be functions of the form (polynomial$\times e^{-x^2}$).

Example 8.1. An "algebraic" construction of Lebesgue measure on $\mathbb{R}$. Our algebraic measure will just be a linear map from (polynomials$\times e^{-x^2}$) to $\mathbb{R}$. Lebesgue measure is supposed to be translation invariant; what does this mean for us? Polynomials times guassians are not invariant under translation, but they are invariant under *infinitesimall* translation (i.e. differentiation). By the way, if you're used to thinking of Lie algebras and Lie groups as the same thing, that doesn't work in infinite dimensions. Here we have an action of a Lie algebra on an infinite-dimensional space which doesn't integrate to an action of the group. So translation invariance means that $\int \frac{d}{dx}(poly \times e^{-x^2} dx = 0$. So we want to find a linear map from (poly$\times e^{-x^2}$)$/\frac{d}{dx}(poly \times e^{-x^2})$ to $\mathbb{R}$, which is easy because this space is 1-dimensional (modulo derivatives, everything is a multiple of e^{-x^2}. This is more or less equivalent to saying that there is a unique translation invariant measure up to scaling. ◇

Remark 8.2. There are some minor advantages to this algebraic approach. It works for $e^{Q(x)}$ where Q is a nonsingular quadratic form in n-dimensional space (it doesn't have to be positive definite). For example, you can define $\int poly \times e^{x^2} dx$ algebraically. ◇

In finite dimensions, everything works just like in the 1-dimensional case (there is a unique translation invariant measure). What about in infinite dimensions? It turns out that the space (poly$\times e^{-x^2}$)$/\frac{d}{dx}(poly \times e^{-x^2})$ is usually infinite-dimensional. So instead of just one candidate for a gaussian measure, we have an infinite dimensional space of candidates. This turns out to the fact that Feynman diagrams are defined up to an infinite-dimensional space of ambiguities.

Remark 8.3. There is a major disadvantage of this algebraic approach. We want to integrate (say) $e^{-\lambda \int \phi(x)^4 dx}$. In the analytic approach, this failed because ϕ^4 blew up almot everywhere. In the algebraic approach, this fails because $\exp(-\lambda \int \phi^4 dx)$ is not a polynomial in ϕ. But we can expand it as a formal power series, and each term can be defined. So we can come up with a formal power series as the answer. Nobody knows how to get around this problem. The answer you get is always a formal power series which usually doesn't converge. ◇

How do we define $\int (\text{poly}(\phi)) e^{-\text{quadratic}(\phi)} D\phi$, where the quadratic thing will usually look like $\int m^2\phi^2 + (\sum \partial^2\phi)\phi dx$ and the polynomial term will look like $\int \phi^4 dx$ times similar terms. Formally we just copy what happens in the finite dimesnional case: the integral can be written as a (finite) sym of Feynman diagrams. The propagator for these diagrams, $\Delta(x_1, x_2)$, is given by the inverse of $m^2 + \partial^2$, which is more or less the Greens function. [[★★★ $\int \phi(x_1)\phi(x_2) e^{-\int m^2\phi^2 + (\partial^2\phi)\phi dx} D\phi$]]

First attempt at the definition: the integral is given by the usual sum of Feynman diagrams. Suppose we want to compute $\int \phi(x_1)^4 \phi(x_2)^6 \phi(x_3)^4 e^{-\int m^2\phi^2 + (\partial^2\phi)\phi dx} D\phi$. then we take the sum over all ways of joining up dots with valance four, six, and four. This leads to a distribution in x_1, x_2, and x_3.

Another way to think of a distribution is to plug in a test function and integrate it. So we can think of it as $\int \int \phi(x_1)^4 f_1(x) dx \times \int \phi(x_2)^6 f_2(x) dx \cdots e^{-\int \ldots dx} D\phi$ (real valued) where f_1, f_2, f_3 are smooth compactly supported functions. These things are polynomials in $\phi(x)$ and derivatives, where x is a smooth function. This is formally the same thing as a Lagrangian (density). So this can be thought of as a linear map from the symmetric algebra of actions of compact support to $\mathbb{R}$. So that's what a guassian measure is, except that this whole thing doesn't actually work. This definition fails because the product of distributions is not defined (because of ultraviolet divergences). However, it is *almost*

well-defined. The product of propagators is well-defined except on the diagonal of $(\mathbb{R}^n)^{\#pts}$. This small ambiguity is controlled by *renormalization.*

The plan of attack is as follows. (1) Define an infinite dimensional space of possible gaussian measures as linear maps Sym(actions) $\to \mathbb{R}$ satisfying some conditions. (2) find a group of renormalizations (acting on Sym(actions)) acting simply transitively on the space of gaussian meassures.

We definitely do not have a canonical Gaussian measure (there are obstructions, called anomalies, proving that you can't get such a thing). However, any two Gaussian measures are equivalent, meaning that there is a unique renormalization taking one to the other. This isn't really anything new.

Example 8.4. A translation invariant measure on a finite-dimensional real space is *not* unique (because you can multiply by a constant). However, the group of positive reals acts simply transitively on the space of such measures. ◇

The difference between the finite-dimensional case and the infinite-dimensional case is that in the finite-dimensional case, we have a 1-dimensional abelian group, and in the infinite-dimensional case we get an infinite-dimensional non-abelian group.

9 RB 10-23

We're going to continue trying to figure out what a gaussian Feynman measure. Goals:

1. *Define* what we mean by a Gaussina Feynman measure.
2. Define a group of renormalizations that acts simply transitively on Feynman measures.

The setup: 1 hermitian scalar field ϕ on (say) Minkowski space $M = \mathbb{R}^{1,d-1}$. We assume we're given a well-behaved propagator Δ that is a distribution on $M \times M$. Most of the constructions of Feynman measure really only depend on the choice of this Δ.

Recall that a Feynman measure is a map from Sym^*(compactly supported actions) to $\mathbb{R}$. A compactly supported action is something that looks roughly like $\int f(x)\phi(x)^4\,dx$ where f is a smooth compactly supported function. The reason for making it compactly supported is to eliminate infra red divergences. The result is not translation invariant, but we'll see how to get that back. This $\phi(x)^4$ can be any polynomial in derivatives of ϕ. PT: you don't have a preferred action corresponding to the propagator? RB: no; it's a little misleading to call them actions, actually. It looks like you're looking at the space $C_0^\infty M \otimes$ (polys in ϕ and derivatives), but you have to mod out by the images of derivatives ∂_i to get actions. You can do

$$\int \underbrace{\int f\phi^4 dx \int f\phi^6 dx}_{\in \mathrm{Sym}^*(\text{actions})} e^{i\int m^2\phi^2 + \phi^2\partial^2\phi dx} \mathcal{D}\phi \to \mathbb{R}$$

We require that a Feynman measure have some property. Formally, this integral can be written as a sum of Feynman diagrams (products of propagators, well-defined up to distribution on the diagonal).

Definition 9.1 (First)**.** A *Feynman measure* is a linear function $\mathrm{Sym}^*(\text{actions}) \to \mathbb{R}$ which can be obtained by summing over Feynman diagrams satisfying the conditions mentioned earlier (that for any edge between x and y, you can remove the edge at the cost of adding a factor of $\Delta(x, y)$). ◇

It would be nice to have a definition which doesn't explicitly talk about summing over Feynman diagrams.

Definition 9.2 (Second)**.** A *Feynman measure* is a linear function $\int\colon \mathrm{Sym}^*(\text{actions}) \to \mathbb{R}$ with the following property. If $a = \int \phi^3 \int \phi^4 \int \phi^2$ and $b = \int \phi^4 \int \phi^6$ are in $\mathrm{Sym}^*(\text{actions})$ *have disjoint supports* (there are hidden functions with compact support, which I'm too lazy to write), then $\int ab$ is

$$\sum \int a' \times \prod(\Delta) \times \int b'$$

summing over all ways to joing a ϕ in a to a ϕ in b, where b' is b without the ϕ's joined to a (and a' similar), and where $\prod(\Delta)$ is a product over propagators where something in a is joined to something in b. ◇

Whichever definition you like, the result is an infinite-dimensional space of Gaussian Feynman measures. Instead of trying to find a canonical element of this space (which is not possible in general), we try to find a group acting transitively on it.

Construction of group of ("finite") renormalizations

What should this group look like?

(1) First of all, it is a subgroup of $GL(\mathrm{Sym}^*(\text{actions}))$, invertible linear maps $\mathrm{Sym}^*(\text{actions}) \to \mathrm{Sym}^*(\text{actions})$ (which don't have to preserve the grading). This acts (by definition) on $\mathrm{Sym}^*(\text{actions})$ and on its dual (which contains Feynman measures). By the way, the reason for using these huge spaces is that the formulas become more transparent; we'll see an example of this in a moment.

(2) We need renormalizations to action on Lagrangians. The reason is that we want a map from Lagrangians×measures to quantum field theories. In physics books, they pretend like there is only one measure, but this is false. Both Lagrangians and measures are acted on by renormalizations, and we want the map to QFTs to be invariant under the renormalization action. Lagrangians are more or less the same as actions, and we've already got an action on actions, but there is something very tricky. The renormalization action on Lagrangians is NONLINEAR. This is one of the reasons renormalization is so hard to understand. The reason it's nonlinear: if you look inside a Feynman integral, it contains a factor of $e^{i\mathcal{L}}$, and it turns out that the action of the renormalization group on these things is linear. $e^{i\mathcal{L}}$ is more or less in $\mathrm{Sym}^*(\text{actions})$ (ignoring convergence problems). We have that the Lagrangians are mapped into $\mathrm{Sym}^*(\text{actions})$ by the exponential map, making Lagrangians a subspace of actions. This induces the action of the renormalization group on Lagrangians.

There is a problem, because there is no reason the action on $\mathrm{Sym}^*(\text{actions})$ should preserve things of the form $e^{i\mathcal{L}}$. Problem: suppose V (which will be the space of actions) is a module over a $\mathbb{Q}$-algebra, and suppose a group G acts on $\mathrm{Sym}^* V$. When does G preserve the subset of elements of the form $e^{\lambda v} \in \mathrm{Sym}^*(V)$ (where λ is infinitesimal, meaning nilpotent). Solution: pretend V is an abelian Lie algebra (you can ignore this if you like), so $\mathrm{Sym}^* V$ is the universal enveloping algebra of V (again you can ignore this if you like), which is, in particular, a Hopf algebra with *coproduct* given by $\Delta(v) = v\otimes 1 + 1 \otimes v$ and extended to make it an algebra homomorphism. A coproduct is a map $\Delta\colon \mathrm{Sym}^* V \to \mathrm{Sym}^* V \otimes \mathrm{Sym}^* V$. The coproduct roughly tell you how something in $\mathrm{Sym}^* V$ acts on the tensor product of two things. Now we can define two special sorts of elements.

Definition 9.3. $g \in \mathrm{Sym}^* V$ is *primative* (Lie-algebra-like) if $\delta(g) = 1 \otimes g + g \otimes 1$. $g \in \mathrm{Sym}^* V$ is *group-like* if $\delta(g) = g \otimes g$. ◇

For example, the primitive elements of $\mathrm{Sym}^* V$ are exactly the elements of V. If G is a group, then the group-ring $\mathbb{C}[G]$ is a Hopf algebra with $\delta(g) = g \otimes g$ for $g \in G$ and the group-like elements of $\mathbb{C}[G]$ can be identified with G.

If g is primative and nilpotent, then $\exp(g)$ is group-like. Conversely, if g is group-like and unipotent (1 plus something nilpotent), then $\log(g)$ is primative.

So elements of the form e^L (if we ignore convergence) are exactly the group-like elements of $\mathrm{Sym}^*(V)$. If we go back to thinking of V as the space of actions, then we want a renormalization group action which preserves the group-like elements of $\mathrm{Sym}^* V$. How can we ensure that an endomorphism preserves the set of group-like elements? An obvious way to do it is to require it to preserve the coproduct $\Delta\colon \mathrm{Sym}^* V \to \mathrm{Sym}^* V \otimes \mathrm{Sym}^* V$.

What on earth are the coproduct preserving maps? Well, a good way to think about this is to dualize everything. What are maps $\mathrm{Sym}^* V \to \mathrm{Sym}^* V$ preserving the product $\mathrm{Sym}^* V \otimes \mathrm{Sym}^* V \to \mathrm{Sym}^* V$. These are easy to identify; they are the same as linear maps $V \to \mathrm{Sym}^* V$, because any such map can be uniquely extended to an algebra homomorphism. Dualizing, you find that maps preserving coproduct can be identified with linear maps $\mathrm{Sym}^* V \to V$.

We've cut down our space a bit, but this space of linear maps is still too big, so we need some more conditions. You may be thinking, "why not require that the product is preserved as well?" Well, then it would be too small to act transitively on all measures.

(3) Renormalizations ALMOST preserve products: $g(ab) = g(a)g(b)$ if a and b have disjoint support.

(4) Renormalizations fix $1 \in$actions (not the same thing as $1 \in \mathrm{Sym}^*(\text{actions})$).

(5) Renormalizations commute with the group of sections of the vector bundle which I'll explain next week.

This more or less defines what a renormalization is. It is clear that renormalizations form a group. This group acts transitively on measures (I'll explain this next week).

10 RB 10-30

Last time I was explaining Gaussian Feynman measure, which is formally something like $e^{i\text{quadratic Lagrangian}}\mathcal{D}\phi$. This is more or less a linear map from the symmetric algebra of compactly supported actions to $\mathbb{R}$. Remember that a compactly supported action is something which looks like $\int f(x)\phi(x)^4\,dx$ where f is compactly supported.

A *renormalization* is a linear map $\mathrm{Sym}^*(\text{actions}) \to \mathrm{Sym}^*(\text{actions})$ preserving certain structures:

1. Renormalizations preserve the coproducts Δ of $\mathrm{Sym}^*(\text{actions})$. Recall that $\mathrm{Sym}^* V$ always has a coproduct $\Delta(v) = v\otimes 1 + 1\otimes v$ extended as an algebra homomorphism. The reason we want renormalizations to preserve Δ is that we want renormalizations to act on the set of things of the form $e^{(\text{action})}$. Remember that this is how we got a (nonlinear) action of the group of renormalizations on the space of actions. PT: do you have the renormalization act on the quadratic Lagrangian or do you fix it once and for all? RB: normally you split the Lagrangian into a quadratic bit and the self-interaction part, and this splitting is non-canonical. Depending on how you split it, the answer might be yes or no.

2. Renormalizations almost preserve the product $\mathrm{Sym}^*(\text{actions}) \times \mathrm{Sym}^*(\text{actions}) \to \mathrm{Sym}^*(\text{actions})$. More precisely, $g(ab) = g(a)g(b)$ whenever a and b have disjoint supports. The reason we want this is to get an action on Feynman measures.

3. Renormalizations commute with the action of sections of the vector bundle ϕ (whose sections are fields) on actions. This is a boring condition. The reason for putting it in is that it cuts down the size of the group of renormalizations so that it acts *simply* transitively on Feynman measures. If you don't put it in, everything works, but your group is too big. BD: "Feynman measure" is just associated to the quadratic part of the Lagrangian? RB: yes. For the non-quadratic part, you have to expand as a power series.

4. Renormalizations preserve 1 in the space of actions. Again this is a boring condition.

The definition of the group of renormalizations looks really technical and hairy, but all you care about is that it acts transitively on measures.

There is a problem here: writing down explicit renormalizations is rather hard. There is an easy construction of explicit renormalizations: take exp(infinitesimal renormalization). To write down a renormalization, you have to know what it does on all Feynman diagrams, which is a pain, but an infinitesimal renormalization only has to be specified on Feynman diagrams of a given order (which you can make zero on all but one of them). An infinitesimal renormalization is something which satisfies infinitesimal versions of all the conditions above.

Now I want to give a vague sketch of why renormalizations act transitively on Gaussian Feynman measures. A proper proof of this would require many technical details which are boring. Suppose M_1 and M_2 are two Feynman measures, and suppose they differ on some Feynman diagram F, but are the same on all smaller Feynman diagrams (and all others of the same size). The idea is to find a renormalization g which fixes all smaller Feynman diagrams such that $g_1(M_1)(F) = M_2(F)$. If we can do this, then we can repeat this an infinite number of times to get the measures to agree on all Feynman diagrams (you can check that the renormalizations converge because there are only finitely many diagrams of a given size).

What is the difference between $M_1(F)$ and $M_2(F)$. Since M_1 and M_2 are the same on all smaller diagrams, the difference $M_1(F) - M_2(F)$ is a distribution with support on some diagonal (remember that the value of a measure on a Feynman diagram is determined by the value on smaller diagrams up to some distribution with very limited support). PT: How are you thinking of these Feynman diagrams as elements of $\mathrm{Sym}^*(\text{actions})$. RB: you can think of measures as functions $\mathrm{Sym}^*(\text{actions}) \to \mathbb{R}$, but you can also think of it as a map $\mathrm{Sym}^*(\text{polys in } \phi \text{ and derivatives}) \to (\text{distributions})$, which you can expand in terms of Feynman diagrams, though I really didn't mean to do that. Since I'm only giving a vague sketch, I'll pretend like this problem doesn't exist.

We're going to construct this renormalization as something of the form $e^{\text{infinitesimal renormalization}}$. The exponential roughly makes sure that the result is well behaved on larger diagrams. So we need to find an infinitesimal renormalization g with $g(M_1)(F) = M_2(F)$. Notice the following things:

1. g is a map $\mathrm{Sym}^*(\text{actions}) \to \mathrm{Sym}^*(\text{actions})$.
2. g is determined by $\mathrm{Sym}^*(\text{actions}) \to (\text{actions})$ because it respects the coproduct and actions generate $\mathrm{Sym}^*(\text{actions})$.
3. Using that g commutes with sections of bundles, we can reduce to a map $\mathrm{Sym}^*(\text{actions}) \to \mathbb{R}$. These are related to distributions on products of space-time.
4. g almost preserves products, which implies that these distributions have support on the diagonal.

So unravelling the definitions, we see that a choice of an infinitesimal renormalization at each step more or less corresponds to the ambiguity, a distribution with support on the Feynman diagram corresponding to the difference between two Gaussian Feynman measures. So we have just enough freedom in choice of g to make $g(M_1)(F) = M_2(F)$. As I said earlier, this is really all you care about when you think about the group of renomalizations.

Actually, there is another problem: I haven't actually constructed a single Gaussian Feynman measure. I've vaguely sketched that given any two measures, there is a renormalization taking one to the other. How can we show that there is at least one Gaussian Feynman measure? There are two proofs, one of which is an existence proof, which I'll vaguely sketch, and the other honestly constructs the measure.

For each Feynman diagram, we need to extend a distribution on $M^n \smallsetminus \text{diagonal}$ to M^n. If we use translation invariance, this reduces to the following basic problem in distribution theory. If we have a distribution on $\mathbb{R}^k \smallsetminus \{0\}$, can we extend it to all of $\mathbb{R}^k$? The answer in general is no. A typical example is e^{1/x^2}. This is so big near the origin that there is no way to extend it to a distribution, though it can be extended as a hyperfunction. Hyperfunctions are extensions of distribution with the property that funtions can always be extended to hyperfunctions. We don't want to use hyperfunctions because they give me a headache. A better answer is that you usually can extend distributions unless you've built it so that you can't extend it. Any distribution with "mild growth" near 0 can be extended. In practice, any naturally occuring distribution will be good enough. In particular, if the original propagator is "reasonable", then all distributions can be extended.

I won't go into this too much because I want to construct an explicit example of a Feynman measure. In order to do this, I want to discuss extension of distributions in detail.

Example 10.1. Take $f(x) = \frac{1}{|x|}$ on $\mathbb{R} \smallsetminus \{0\}$. This can be extended to $\mathbb{R}$, but there is no canonical way of doing it. For example, $\frac{1}{|x|} = \frac{d}{dx}(\log x \cdot \sin x)$ for $x \neq 0$. This is a locally integrable function, so it is a distribution. If you have any distribution, its derivative is another distribution. This looks perfectly canonical, but there is a catch. $f(x)$ is homogenous of degree -1 (that is, it is invariant under some sort of rescalings of the reals), and we would like the extension to also be of degree -1. If there were a unique way to extend, it would automatically be of degree -1. There is no way to do this. The problem is that $\log|x|$ is "not quite of degree 0". Rescaling adds a constant to $\log|x|$, which messes up rescaling of the extension:

$$\frac{d}{dx}\big((\log x + c)\sin x\big) = \frac{d}{dx}\log(x)\sin(x) + c\frac{d}{dx}\sin(x)$$

So when we rescale we pick up delta functions at the origin. You can see that we're not doing something stupid as follows. All possible extensions of $\frac{1}{|x|}$ differ by a distribution with support at the origin. Such distributions are just spanned by deriviatives of $\delta(x)$. Since the higher derivatives $\big(\frac{d}{dx}\big)^n\delta(x)$ of $\delta(x)$ are homogeneous of degree $-1-n$, we should only be thinking about $\delta(x)$. Degree n means $x\frac{d}{dx}f = nf$, or $\big(x\frac{d}{dx} - n\big)f = 0$. It turns out that for $f(x) = \frac{d}{dx}\log(x)\sin(x)$, $\big(x\frac{d}{dx}+1\big)^2 f = 0$, so it is some kind of generalized eigenvector. All extensions that are generalized of degree -1 are given by $\frac{d}{dx}\log(x)\sin(x) + c\delta(x)$. Rescaling acts *transitively* on these things by $c \mapsto c + \log(\lambda)$. So there are no homogenous extensions of degree -1. $\diamond$

Now let's consider $\frac{1}{|x|^s}$ as a meromorphic-distribution-valued function of s. The idea is going to be that if you send s to 1, you find that there is a pole, and this pole measures the problem. We'll do that next week.

11 RB 11-06

Last week we were looking at the following example. $\frac{1}{|x|}$ is a distribution for $x \neq 0$. It can be extended to a distribution for $x = 0$, but not in a canonical way. This is supposed to be a toy example of what is going to happen for Feynman diagrams. First of all, let's look at it in more detail.

Here is an idea for finding a canonical extension: look at $|x|^s$, where $s \in \mathbb{C}$. If $Re(s) > 0$, this is actually continuous, so it gives a distribution. The idea is to analytically continue it to all values of s. What do we mean by that? We think of $|x|^s$ as a function of s taking values in the space of distributions in x. This always gives people a bit of a headache when they first come across the concept. The idea is that for each test function $f(x)$, we look at $\int f(x)|x|^s\,dx$, and think of this as an analytic function of s. Analytically continuing is the same thing as saying you can extend this integral for other values of s.

How do we continue $|x|^s$ to $Re(s) \leq 0$. Here's an idea for doing it: differentiate to get $\frac{d}{ds}|x|^s = s|x|^{s-1}\mathrm{sign}(x)$ for $x \neq 0$, $s > 1$. Thus, we get $|x|^{s-1} = \frac{1}{s}\frac{d}{ds}|x|^s\mathrm{sign}(x)$. This $\mathrm{sign}(x)$ is a headache because it is discontinuous. To get rid of it, we differentiate again to get $\frac{d^2}{dx^2}|x|^s = s(s-1)|x|^{s-2}$ for $Re(s) \gg 0$. So we get $|x|^{s-2} = \frac{1}{s(s-1)}\frac{d^2}{dx^2}|x|^s$, giving an analytic extension of $|x|^s$ to a meromorphic function of s (meromorphic because we may get poles at $s = 0, 1$).

Where are the poles and what are the residues (the residue at a pole is a distribution, not a number). For example, we have $s|x|^{s-1} = \frac{d}{dx}|x|^s\mathrm{sign}(x)$. If we try putting $s = 0$, we see that this right hand side is the residue of $|x|^{s-1}$ at $s = 0$. We can figure out the value by just setting $s = 0$, getting $\frac{d}{dx}\mathrm{sign}(x) = 2\delta(x)$. So the residue is a constant times the Dirac delta.

Now let's go back and try to define $\frac{1}{|x|}$ as a distribution. First attempt: take the value of $|x|^s$ at $s = -1$, which fails because there is a pole at $s = -1$ (we've just worked out the residue at this pole). Second attempt: take the constant term of $|x|^s = a_{-1}(x)s^{-1} + a_0(x) + a_1(x)s + \cdots$ (where the a_i are distributions). This doesn't work because of the following subtle problem. The constant term of a meromorphic function at a singular point is not canonical (for example, it changes with a change of variables). Suppose I take $f(s) = s^{-1} + b + \cdots$, and take $s = t + ct^2 + \cdots$, with

$s^{-1} = t^{-1} - c$. Then $f(t) = at^{-1} + b$+constant depending on a plus more stuff. So the constant term is only canonical up to addition of a multiple of the residue. If we were working with $\mathbb{C}$-valued functions, this would be useless, but since we're working with distribution-valued functions, this gives is some information.

Similarly, if $f(s) = a_{-n}s^{-n} + \cdots + a_0$, then the constant term changes by multiples of a_{-n}, ..., a_{-1}. So we can't pick out the constant term, but we can pick out a subspace of distributions which could be constant terms. So we have a canonical *family* of extensions of $\frac{1}{|x|}$ to $x = 0$ by taking $\frac{1}{|x|}$ =constant term of $|x|^s$ at $s = -1$ plus a complex nmber times the residue at $s = -1$. The residue tells us the ambiguity in extending $|x|^{-1}$. If we have a pole of order k, then we find that there is a k-dimensional space of extensions, controlled by the singularity. $|x|^s$ has poles at $s = -1, -3, -5, \ldots$.

Example 11.1. Consider $\int_{-\infty}^{\infty} |x|^{s-1}e^{-x^2}\,dx$ (think of e^{-x^2} as a test function). If we change x to x^2, this is $2\int_0^{\infty} e^{-x}|x|^{s/2-1}\,dx = 2\Gamma(s/2)$. We know that Γ has an analytic continuation, and lots of poles all over the place. Since $|x|^{s-1}$ is meromorphic at $s = 0, -2, -4, \ldots$, we see that $\Gamma(s)$ is meromorphic at $s = 0, -1, -2, \ldots$. So the meromorphic continuation of $|x|^s$ is more or less the meromorphic continuation of Γ. If you unwind this, you find that the proof of meromorphic continuation of $|x|^s$ is more or less the same as the proof of meromorphic continuation of Γ. $\diamond$

Summary: to define $|x|^{-1}$, we first look at $|x|^s$ for s large, analytically continue it to $s = -1$, take the constant term plus multiples of the singular parts.

Now let's consider the more complicated case of Feynman diagrams.

Problem: Feynman diagram is a product of propagators $\Delta(x_i - x_j)$, and the product is sometimes not defined if $x_i = x_j$. So we have the same sort of problem as for $|x|^s$.

Strategy:

1. Replace the propagator by some holomorphic family of propagators $\Delta_s(x_i - x_j)$ for $s \in \mathbb{C}$ so that if $Re(s) \gg 0$, then $\Delta_s(x_i - x_j)$ is continuous (say) with $\Delta_0 = \Delta$.
2. For large s, the product of propagators is defined, so we get a function from s to distributions
3. Try to take some analytic continuation in s to $s = 0$. We usually find that it is meromorphic at $s = 0$ with a pole of finite order. This gives us a reasonable space of extensions whose dimension is the order of the pole. The Feynman diagram is given by the constant term plus multiples of the singular parts at $s = 0$.

(4. Try to eliminate poles by applying a renormalization.)

How do we do step (1)? What should we use as $\Delta_s(x)$? The propagator is a bit of a pain in x coordinates. Notice that $\Delta(x)$ is the Fourier transform of $\frac{1}{p^2+m^2}$ (suppose we're doing scalar field theory in Euclidean space). Then $\Delta(x)$ has singularities at $x = 0$. Singularities are caused by the fact that $\frac{1}{p^2+m^2}$ does not decrease fast enough at $p = \infty$ (in reasonably high dimensions). So we should make $(p^2 + m^2)^{-1}$ decrease faster. You can immediately see lots of ways of doing this. Here are a few of the more popular ways.

1. $(p^2 + m^2)^{-1-s}$. This is similar to dimensional regularization. Physicists say that this is "changing the dimension of spacetime to $4 + s$". This has the additional advantage that it actually makes some sense. This still doesn't decrease all that fast as $p \to \infty$.
2. $(p^2 + m^2)^{-1}e^{-(p^2+m^2)}$ (there should be an s in there somewhere), which has several advantages over the previous one.
3. $\int_0^{\infty} \theta(t)e^{-(p^2+m^2)t}\,dt$.

Let's just use the first one. You might think that this gives you a canonical extension. The trouble is that $(p^2 + m^2)^{-1-s}$ is not scale invariant (since it doesn't have the same degree as $(p^2 + m^2)^{-1}$). You can add an extra mass μ, and take $(p^2 + m^2)^{-1}\left(\frac{\mu^2}{p^2+m^2}\right)^s$. This extra mass really does have to show up ... it's not just that you've done something stupid.

Bruce: what exactly does scale invariance mean if $m \neq 0$? RB: it doesn't make a whole lot of sense; I guess you rescale the mass an your spacetime.

Ok, so now we've chosen a family of propagators and we find that (2) works.

Now what about step (3)? This involves a certain amount of work. For simplicity, we take $m = 0$. Then the propagator is $(x^2)^{1-d/2}$ times some constant. Adding in a factor of s means we have a family of propagators $(x^2)^s$ (let's ignore the $d/2$ part for now). So a Feynman diagram like [[★★★ picture]] gives us something of the form (polynomial)^s. This is defined for $Re(s) > 0$ and we want to analytically continue it.

General problem: given a polynomial p (say positive), can we analytically continue $p(x)^s$ to all complex s? The answer is YES! This is a theorem of Bernstein and Sato (at about the same time).

To see how to do this, look at the case of $p(x)^s = x^s$. Consider $\frac{d}{dx}x^s = sx^{s-1}$. Then we have $sp(x)^{s-1} = \frac{d}{dx}p(x)^s$. We can write this as $b(s)p(x)^{s-1} = (\text{diff operator with polynomial coeffs})p(x)^s$. This $b(s)$ is called the *Bernstein polynomial of* $p(x)$. If we can find such a relation, then we can continue $p(s)$ to all complex s.

Next week will be all about Bernstein polynomials and how you can prove they exist.

12 RB 11-13 Bernstein polynomials

In the last lecture we ran into the following problem. We wanted to define a Feynman diagram as something like a product of propagators $\prod \Delta_s(x_i - x_j)$, where $\Delta_s(x_i, x_j)$ should be analytic in s and continuous for $s \gg 0$ (so the product is well defined) and we analytically continue the product to the s we're interested in (where there actually turns out to be a pole).

More generally, suppose $p(x)$ is any polynomial in $x = (x_1, \ldots, x_n)$. Can we analytically continue $|p(x)|^s$? The answer is yes; you can use a Bernstein polynomial $b(s)$, which has the property that $b(x)p(x)^{s-1} = (\text{some polynomial diff operator})p(x)^s$. You pick up poles from the zeros of $b(s)$ (the poles of the analytic continuation will be at integer shifts of the zeros of $b(s)$, since you have to repeat).

Example 12.1. $sx^{s-1} = \frac{d}{dx}x^s$. This is essentially how you prove the analytic continuation of the gamma function. ◇

Example 12.2. $p(x) = x_1^2 + \cdots + x_n^2$. This is what you're interested in if you want to raise Laplacians to various powers. Well, $\frac{d}{dx_1}p(x)^s = 2x_1p(x)^{s-1}$, so $\frac{d^2}{dx_1^2} = 2sp(x)^{s-1} + (2x_1)^2s(s-1)p(x)^{s-2}$, so

$$\sum_i \frac{d^2}{dx_i^2}p^s = (2ns + 4s(s-1))p^{s-1}$$

so the Bernstein polynomial is $b(s) = 2ns + 4s(s-1)$. ◇

Example 12.3. $p(x) = x^2 + y^3$. If you manage to find the Bernstein polynomial of this, I'll be very impressed. It takes a huge amount of calculation. The answer $b(s) = (s+1)(s+5/6)(s+7/6)$. Direct calculation is very hard. ◇

The point is that finding Bernstein polynomials in general is very difficult. Fortunately, we don't always need to know $b(s)$ explicitly. Bernstein showed how to prove that it exists without actually producing it. The existence result is actually very powerful.

Application to show the strength of the Bernstein polyomial: the Malgrange-Ehrenpreis Theorem. This is the fundamental theorem of differential operators with constant coefficients. People spent about a hundred years trying to prove it. It says that partial differential operators with constant coefficients have a Green function, i.e. $p(b)f = \delta$ has a solution f for any polynomial p in $\frac{\partial}{\partial x_i}$ with constant coefficients. This is actually false if you allow polynomial coefficients, which was surprising.

Proof. Take Fourier transforms. We have to solve $\tilde{P}(p)\tilde{f} = 1$, where $P(p)$ is a polynomial in P. So we have to find a *distributional* inverse of a polynomial $P(p)$. In the case where it is positive, this is trivial because you can invert it. We can assume $P \geq 0$ by replacing it with $P\overline{P}$ (if we can invert this, it will be easy to invert P). The case $P > 0$ is trivial. You might think that you can then take a limit, but there is a problem. The problem is the complexity of a singularity of $P(p)$. Hironaka's resolution of singularities can be used to do this. The nice thing about Bernstein's polynomial is that it is easier to prove and can often be used as a substitute for Hironaka's resolution of singularities.

Look at $P(p)^{s-1} = a_{-n}s^{-n} + \cdots + a_0 + \cdots$ near $s = 0$. We're using Bernstein to analytically continue and look at the Laurent expansion. $P(p)P(p)^{s-1} = P(p)^s = 1 + (\quad)$ near $s = 0$ (if $P \not\equiv 0$). So $P(p)a_{-n} = 0$, $P(p)a_{1-n}, \ldots$, $P(p)a_0 = 1$, so a_0 is a distributional inverse to $P(p)$ as desired. □

Warning 12.4. The inverse of $P(p)$ is NOT unique. This seems odd because it is easy to prove that it is unique: suppose $aP = 1 = bP$, then $a = aPb = b$. This proof is wrong because multiplication of distributions by polynomials is NOT ASSOCIATIVE! For example, $\left(\frac{1}{x} \cdot x\right) \cdot \delta(x) = \delta(x) \neq 0 = \frac{1}{x} \cdot 0 = \frac{1}{x} \cdot \big(x \cdot \delta(x)\big)$. ⌟

PT: is there a finite-dimensional space of inverses? RB: the space of inverses is usually infinite-dimensional. If you have a function of more than one variable, say $p(x, y)$, then if you have a function on $p(x, y) = 0$, you get something. Even when $p(x, y) = 0$ is a zero-dimensional variety, you can still run into trouble.

Example 12.5. Say you want to solve $\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\right)f(x, y) = \delta$, with f rotationally invariant. The solution is $f(x, y) = \text{const} \cdot \log(x^2+y^2) + \text{arbitrary}$ constant. This arbitrary constant is exactly arising from the pole. ⋄

The point is that this problem of having unknown functions really does turn up in simple natural examples. This shows that when you try to quantize something something, then dialation invariance is automatically broken.

The proof of Bernstein polynomial has lots of useful ideas which should be in every mathematician's kit. It uses $\mathcal{D}$-modules (in fact, a lot of $\mathcal{D}$-module stuff was invented to prove this). A $\mathcal{D}$-module is a module over the ring $\mathcal{D} = k[x_1, \ldots, x_n, \partial_1, \ldots, \partial_n]$. This is a non-commuting ring in the sense that $[x_i, x_j] = 0 = [\partial_i, \partial_j]$, and $\partial_i x_j = x_j\partial_i + \delta_{ij}$. This is basically the universal enveloping algebra of a Heisenberg algebra.

Example 12.6. S: all smooth functions on $\mathbb{R}^n$ with x_i acting by multiplication by x_i and ∂_i acting by $\frac{\partial}{\partial x_i}$. ⋄

Example 12.7. Suppose you have a system of PDEs D_1, D_2,... with polynomial coefficents. Then we get a $\mathcal{D}$-module $M = k[x_i, \partial_i]$/left ideal generated by D_i. Note that Smooth solutions to the system are exactly the same thing as $\mathcal{D}$-module homomorphisms from M to S. So modules over $\mathcal{D}$ should give you some insight into solutions to systems of PDEs. ⋄

Recall some basic commutative algebra. Recall that the *Hilbert polynomial* of a finitely generated graded $k[x_1, \ldots, x_n]$-module $M = \bigoplus_{i\geq 0} M_i$ is given by $p(i) = \dim_k \bigoplus_{j=0}^{i} M_j$ for some polynomial p for large i. This polynomial in some sense measures the size of the module. The existence of the Hilbert polynomial is easy using induction on n. For $n = 0$, this is completely trivial because M is a finite dimensional vector space so $M_i = 0$ for $i \gg 0$, so $p(i) = \dim M$ for $i \gg 0$. For $n > 0$, you consider $M_{i-1} \xrightarrow{\cdot x_n} M_i$ and look at the kernel and cokernel:

$$0 \to \ker(x_n)_{i-1} \to M_{i-1} \xrightarrow{\cdot x_n} M_i \to (M/x_nM)_i \to 0$$

Both the kernel and cokernel are modules over $k[x_1, \ldots, x_{n-1}]$ so by induction, they both have Hilbert polynomials. So $\dim M_i - \dim M_{i-1}$ is a polynomial for $i \gg 0$ by induction. But then $\dim M_i$ is also a polynomial (of one degree higher) for $i \gg 0$.

If a polynomial $p(n)$ is integral for all large integer n, p need not have integer coefficients. For example, $p(n) = n(n+1)/2$. However, $p(n)$ is an

integral linear combination of $\binom{n}{k}x^k$. This is an easy exercise. So if p has degree k, the leading coefficient is $\frac{m}{k!}n^k + \cdots$ for some integer m.

If M is a module with Hilbert polynomial $p(n) = \frac{m}{k!}n^k + \cdots$, then k is called the *dimension* of M and m is called its *multiplicity*. These are the most important measure of size of M. The complete Hilbert polynomial depends on choices, but these values do not.

The ring we want to work with is not commutative so we have the following problem: we want to turn the non-commutative ring $\mathcal{D}$ into a commutative ring so that we can apply the nice theory of Hilbert polynomials to its modules. Well, $\mathcal{D}$ is very close to being commutative. We have the *Bernstein filtration* $B_0 \subseteq B_1 \subseteq \cdots$, where B_i is the set of polynomials in x_i and ∂_i of degree less than or equal to i. You have to be careful talking about degree; there is no such thing as a homogeneous polynomial (e.g. is $x_i\partial_i - \partial_i x_i = 1$ of degree 2 or 1?), but you can talk about thing of degree *at most* i. B_i and B_j commute modulo terms in B_{i+j-1}. So we can form a commutative ring of about the same size: $B_0 \oplus B_1/B_0 \oplus B_2/B_1 \oplus \cdots$. This has the same generators, but is now a commutative polynomial ring.

Now let's take a finitely generated module M over this non-commutative ring $\mathcal{D}$, generated by some finite vector space M_0. Let $M_i = B_i \cdot M_0$, so we have $M_0 \subseteq M_1 \subseteq \cdots$. Notice that $B_iM_j \subseteq M_{i+j}$, so $M_0 \oplus M_1/M_0 \oplus M_2/M_1 \oplus \cdots$ is a graded module over $\bigoplus B_i/B_{i-1}$, which is about the same size as the original module. Now we can start applying Hilbert polynomials and things.

Warning 12.8. This module $\bigoplus M_i/M_{i-1}$ depends on the choice of M_0. The Hilbert polynomial changes if M_0 changes. So we can't talk about the Hilbert polynomial, but it turns out that the multiplicity and dimension of $\bigoplus M_i/M_{i-1}$ do not depend on the choice of M_0, so we can talk about the multiplicity and dimension of M. ⌟

Next week we'll finish off the proof of existence of the Bernstein polynomial.

13 RB 11-20

[[★★★ Using absent-mindedness, I missed this class.]]

14 RB 11-27

Today's lecture will be on the following question: how do you reconstruct a QFT from Feynman diagrams? First I'll explain what the problem is. We've more or less shown how to define Feynman diagrams. For each diagram, we've shown how to construct some kind of distribution (in a way that depends on the choice of Gaussian Feynman measure which is quite complicated to construct). Assume we've solved the problems of constructing the Feynman measure. Then we can construct Green's functions, defined as

$$\int \phi(f_1)\phi(f_2)\cdots e^{iL(\phi)}\mathcal{D}\phi$$

Where L is a Lagrangian of the form (free part) + λ(integral part). This is a formal power series in λ with coefficients which are distributions on M^n.

But a QFT isn't a bunch of distributions. A QFT consists of a space H (which is something like a Hilbert space) which is a module over $\mathbb{C}[[\lambda]]$ with a sesquilinear pairing $(,)$, with operators $\phi(f)$ for every smooth funcion f of compact support. These $\phi(f)$ should satisfy some axioms, which we listed earlier. Once you construct H and $\phi(f)$, it usually isn't too bad to check the axioms. So how do we go from a pile of Green's functions to a Hilbert space with some operators? First of all, we use the GNS construction, or the Wightman reconstruction theorem. Recall that Wightman distributions on M^n are $W(f_1,\ldots,f_n) = \langle vac|\phi(f_1)\cdots\phi(f_n)|vac\rangle$. From these distributions, you can reconstruct the QFT. The Wightman distributions give a "state" (morally, but not really) on the free algebra A generated by all operators $\phi(f_i)$ (i.e. a linear map $\omega\colon A \to \mathbb{C}$, given by $\omega(\phi f_1\ldots,\phi f_n) \mapsto W(f_1,\ldots,f_n)$. Now we can reconstruct H by the GNS construction. We define the inner product on A by $(a,b) = \omega(ab^*)$. Then $H = A/\ker(,)$. So to reconstruct the field theory, it is enough to reconstruct the Wightman distributions.

Feynman diagrams give you loads of distributions (Green's functions) $G(f_1,\ldots,f_n)$. The Wightman reconstruction theorem says that by starting with loads of distributions $W(f_1,\ldots,f_n)$, we can get a QFT. This suggests that we construct the QFT by taking $W(f_1,\ldots,f_n)$ to be $G(f_1,\ldots,f_n)$. This doesn't work for a very simple reason. The Green's functions are obviously symmetric in the f_i. PT: there are two meanings of $\phi(f)$ here. One is an operator on H. What is the other meaning? RB: the integral defining the Green's function is just obtained by integrating the distribution coming from the Feynman diagram against the functions $f_1,\ldots,f_n$. On the other hand, the Wightman distributions are not symmetric because the operators $\phi(f_i)$ do not commute (unless the supports are spacelike separated.

However, the idea is not completely stupid and it nearly works, but needs some modification. If f_1 and f_2 have spacelike separated supports, then the operators $\phi(f_1)$ and $\phi(f_2)$ commute. So the Wightman distributions $W(f_1,\ldots,f_n)$ are symmetric if the supports of the f_i are spacelike separated. In this case, we can identify Wightman distributions with Green's functions. The key point is that if you know the Wightman distributions on spacelike points, this determines the Wightman distributions at all points as follows.

Wightman distributions $W(x_1,\ldots,x_n)$ (pretend they are functions) are in fact boundary values of holomorphic Wightman functions $W(z_1,\ldots,z_n)$, where $z_i \in M\otimes\mathbb{C}$ and $Im(z_j - z_i) \in C$ (positive cone) (see some basic book on QFT, like Streater-Wightman). Moreover, these holomorphic functions can be analytically continued to a larger region, including all totally spacelike points, and we know they have to be Green's functions on totally spacelike points.

Schematically, we have (symmetric) Green's functions, which are distributions on M^n. These give you Wightman functions $W(z_1,\ldots,z_n)$ by analytic continuation from spacelike points (completely ignore what the Green's functions do on non-spacelike points). These Wightman functions are symmetric (because they are analytic continuations of symmetric functions). Taking certain boundary values, you get Wightman distributions. Finally, applying the GNS construction (Wightman reconstruction theorem), we get the QFT.

The Wightman distributions are not symmetric. Why is the boundary value of a symmetric function not symmetric? Because there is more than one way to take a boundary value. For example, suppose you take $\sqrt{x}$, defined for $x \not\leq 0$ on the complex plane. It has two possible boundary values, $\sqrt{x+0i}$ and $\sqrt{x-0i}$. Similarly, for Wightman functions, you can approach from many different directions (this thing is codimension n, so there are many different ways to approach). PT: which one do you

pick? RB: given a point $(x_1, \ldots, x_n)$, approach from $(z_1, \ldots, z_n)$ with $Im(z_i - z_j) \in C$ (positive cone in spacetime tensor $\mathbb{C}$) when $i > j$.

There are several problems with this approach.

1. We need to show that the analytic continuation exists. If you work in Minkowski spacetime, you can go through the mess of writing it out.

2. The analytic continuation doesn't even make sense on curved spacetimes. As far as I know, there is no such thing as the complexification of a Riemannian manifold.

3. Infrared divergences, which I've been brushing under the carpet for most of the semester.

The problem is that Feynman diagrams give distributions, so we can handle things like interactions of the form $\int \lambda(x)\phi(x)^4\,dx$ where λ is smooth of compact support. But we want λ to be 1, which is not of compact support. Interactions usually look like $\lambda \int \phi(x)^4\,dx$. We want it to be 1 so that it is invariant under the action of the Poincaré group. If you made λ of compact support, everything would converge nicely, but your theory wouldn't be Lorentz invariant.

So we want a way of reconstructing the QFT which doesn't involve analytic continuation on the complexification (which may not exist) and also deals with this infrared divergence.

Recall the two cut propagators Δ^+ and Δ^-. The ordinary propagator $\Delta(x)$ is basically the Fourier transform of $((p + i\varepsilon)^2) + m^2)^{-2}$. Δ^+ is the Fourier transform of something with support on $p^2 = m^2$ and Δ^- similar. $\Delta(x) = \Delta^+(x)$ if x is not in the negative cone and $\Delta(x) = \Delta^-(x)$ is x is not in the positive cone. Similarly, we have the Feynman propagator (and its complex conjugate) $\Delta^*(x) = \Delta^-(x)$ if x is not in the negative cone.

Now I want to introduce some slightly more complicated expressions. Formally, we have

$$\int \phi(f_1)\phi(f_2)e^{i\text{quadratic}+\lambda\phi^4}\mathcal{D}\phi = \sum \text{Feynman diagrams}$$

where the propagators on the right are Feynman propagators. Now I want to define a sort of generalization of this.

$$\phi(f_1)\phi(f_2)e^{iL_1} \;\Big|\; \phi(f_3)\phi(f_4)e^{iL_2} \;\Big|\; \cdots$$

These vertical lines mean the following. Sum over Feynman diagrams in the usual way, but any line which passes over one of these vertical lines is going to be a cut propagator instead of a Feynman propagator, unless I put a minus sign (e^{-iL_2}), in which case we use the complex conjugate of the Feynman propagator. PT: to define the Feynman diagrams, you had to make choices with Bernstein polynomials. RB: yes, to define these with just Feynman propagators, there was some ambiguity, but if you multiplying by some cut propagators is actually easy because of their wave front sets.

Now I want to define something which is more or less a scattering matrix with a perturbation by a field. $S(f_1f_2)\overline{S}(f_3f_4)S(f_5f_6)\cdots$ is defined to be this sum of Feynman diagrams with $L_1, L_3, L_5, \ldots$ to be the interaction and $L_2, L_4, \ldots$ to be the complex conjugates. This S should be thought of as a scattering matrix.

The Wightman distributions $W(f_1, \ldots, f_n)$ will turn out to be $S_1(f_1)\overline{S}S(f_2)\overline{S}S(f_3)\cdots$. There is something funny going on because we've introduced these cut propagators and these inverse scattering matrices. You can think of this as scattering with a source f_1, then undo the scattering, then scatter again with source f_2, then undo, and so on. This will actually work for all $x_1, \ldots, x_n$, not just those with $x_i - x_j$ spacelike.

Now we need to check that this definition of $W(f_1, \ldots, f_n)$ satisfies the Wightman axioms. Most of them are fairly trivial (like Lorentz invariance). The main difficulty is to show that these satisfy locallity, which says that $W(f_1, f_2, \ldots, f_n) = W(f_2, f_1, \ldots, f_n)$.

Next week I'll try to explain why this choice of Wightman distributions give you a QFT.

15 RB 12-04

This is the last lecture this semester. For next semester, there is a vague plan to have a seminar on conformal field theory. I'm not really sure what conformal field theory is, and nobody else seems to know either.

Today I want to finish describing this problem of how to get from Green's functions to a quantum field theory. Recall what we did last week. We had these Green's functions given by sums of Feynman diagrams (assuming you can deal with renormalization and regularization). The quantum field theory is given by Wightman distributions. We'd like to take the Wightman distributions to be the Green's functions, but this doesn't work (e.g. Green's functions are symmetric, but Wightman distributions aren't). We take the Wightman distributions to be equal to the Green's functions at points $(x_1, \dots, x_n)$ where $x_i - x_j$ is spacelike.

There is another way of saying this. Wightman distributions can be thought of as analytic continuations of Green's functions. What does this mean for distributions (analytic continuation only makes sense for analytic functions). What is the ananlytic continuation of a distribution? Why does analytic continuation make sense for analytic functions? It is because of the following uniqueness condition: if f is analytic in a connected open set and zero on a non-empty open subset, then f is identically zero. This tells you that analytic continuation is unique if it exists.

We want a similar condition for distributions. Obviously, we can't do this for arbitrary distributions (it isn't even true for smooth functions), so we have to add some condition to the distribution. We could require that the distribution is analytic, but that is too strong. For QFT, the correct condition turns out to be the following. Use *analytic wave front sets*, which are sort of like wave front sets (which tells you where the distribution isn't smooth and the direction in which it isn't smooth). The analytic wave front set tells you where the distribution is not analytic and in what direction it fails to be analytic (so the analytic wave front set is larger than the wave front set). Suppose we choose a proper (doesn't contain a line) closed cone C_x in the cotangent space of each $x \in M$. Then *if f is a distribution with analytic wave front set contained in C_x at each point x and $f = 0$ on some open set, then f is identically zero.* I won't prove this, but I'll give you some examples. You have to be careful because the continuation of f may depend on the choices of C_x. Also, if you start off with a real distribution, its continuation can in general be a complex distribution.

Example 15.1. Let $f = \sqrt{x}$ for $x > 0$ (on $\mathbb{R}$). We want to analytically continue it to a distribution on all of $\mathbb{R}$. There are two ways to do this.

1. Take f to be the boundary value of $\sqrt{x}$ for $Im(x) > 0$.
2. Take f to be the boundary value of $\sqrt{x}$ for $Im(x) > 0$.

The difference is that the analytic wave front sets at zero are opposite cones. ◇

Example 15.2. Suppose $M = \mathbb{R}^n$ and $C_x = C \subseteq (\mathbb{R}^n)^*$ is the same cone at each point. Then the analytic wave front set of f is in C at each point x if and only if (more or less) f is the boundary value of a holomorphic function in the cone $\mathbb{R}^n \oplus i\hat{C}$. Since the cone C is closed and proper, the dual cone $\hat{C}$ has non-empty interior. In general, M may not have a complexification, so it isn't so clear what you mean by saying that f is a boundary value of a holomorphic function on the complexification. ◇

Example 15.3. Any distribution is locally a (finite) sum of boundary values of holomorphic functions. Suppose f is a distribution of compact support on $\mathbb{R}$. Then set $g(z) = \frac{1}{2\pi i}\int_{\mathbb{R}} \frac{1}{z-x} f(x)\,dx$ (this is well defined provided z is not in the support of f.). Then "$f(x) = g(x + i\varepsilon) - g(x - i\varepsilon)$". ◇

The point is that if you select a proper cone at each point, there is a way to make sense of analytic continuation of distributions.

We're trying to find an analytic continuation W of a Green function such that the analytic wave front set is contianed in the region $\{(z_1, \dots, z_n) | z_i - z_j \in \text{cone for } i < j\}$. Last lecture I wrote down a confusing formula for what W is.

$G(x_1, \dots, x_n)$ is a sum of Feynman diagrams that have a node of valence 1 for each x_i and nodes of valence 4 coming from the $\lambda\phi^4$ in the Lagrangian. $W(x_1, \dots, x_n)$ will be similar sums of Feynman diagrams, but with slightly different propagators. Instead of summing over all Feynman diagrams with valence 4 nodes, you sum over diagrams like $x_1 \ \times\times| \times\times| \times\times x_2| \times\times| x_3$, where the x_i are nodes of valence 1 and

the $\times$ are nodes of valence 4. The propagators in the odd buckets are Feynman propagators, in the even buckets are conjugates of Feynman propagators, and the ones that cross buckets are cut propagators.

More generally, let's define $S(f_1, f_2, \dots,)\overline{S}(g_1, \dots)S(h_1, \dots)$ to be a sum of all diagrams as in the previous paragraph, but in the first bucket we allow nodes for the f_i, in the second bucket, we allow some g_i, and so on.

Main properties:

1. $S(f_1, \dots, f_n)\overline{S}(g, h_1, \dots, h_m) = S(f_1, \dots, f_n, g)\overline{S}(h_1, \dots, h_m)$ provided no element of $\operatorname{Supp} g$ is less than or equal to any element of $\operatorname{Supp} f_i$ or $\operatorname{Supp} h_i$.
2. $S\overline{S} = 1 = \overline{S}S$ (i.e. if you don't put any functions, they cancel). This is assuming you've chosen a renormalization already.

Rather than proving these results, let me show you how to use them to prove that the Wightman distributions $W(f_1, f_2, \dots) = S(f_1)\overline{S}S(f_2)\overline{S}\cdots$ satisfy locality.

$$\begin{aligned} W(f_1, f_2) &= S(f_1)\overline{S}S(f_2)\overline{S} \\ &= S(f_1)\overline{S}(f_2)S\overline{S} \\ &= S\overline{S}(f_1, f_2)S\overline{S} && \text{provided } \operatorname{Supp} f_1 \not\leq \operatorname{Supp} f_2 \\ &= S(f_2)\overline{S}(f_1)S\overline{S} && \text{provided } \operatorname{Supp} f_1 \not\geq \operatorname{Supp} f_2 \\ &= W(f_2, f_1) \end{aligned}$$

So we can switch f_1 and f_2 provided their supports are spacelike separated.

We can also handle infrared divergences. I've been saying that these divergences automatically cancel. Suppose you're working with a Lagrangian L with compact support. Add to this another Lagrangian M, also with compact support. I want to find conditions under which addition of M makes no difference to the Wightman distributions. To do this, we can make the following calculation. The new Wightman distributions are given by

$$\begin{aligned} S(f, e^M)\overline{S}(e^{-M})S(f_2, e^M)\overline{S}(e^{-M})\cdots &= S\overline{S}\cdots S\overline{S}S(f, e^M)\overline{S}(e^{-M})S(f_2, e^M)\overline{S}(e^{-M})\cdots \\ &= S(f_1)\overline{S}S\overline{S}\cdots S(e^M)\overline{S}(e^{-M})\cdots \end{aligned}$$

provided no element of $\operatorname{Supp} f_1$ is $\leq$ anything in $\operatorname{Supp} M$. Keep pulling the $S(f_i)$ through and cancel the left over copies of $S(e^M)\overline{S}(e^{-M})$. What is left is the Wightman distribution for the Lagrangian L. So adding M to the Lagrangian does not affect the Wightman distributions provided that no element of $\operatorname{Supp} f_i$ is less than or equal to anything in $\operatorname{Supp} M$.

A similar argiment shows that the same is true if no element of $\operatorname{Supp}\{f_i\}$ is greater than or equal to anything in $\operatorname{Supp} M$.

[[★★★ picture]]

So suppose the support of M lie in the shaded region, then adding M doesn't make a difference. So we're left with a compact region that we have to leave alone. We can mess with the Lagrangian outside of this compact region. Suppose the interaction is given by $\int \lambda(x)\phi(x)^4\,dx$, where $\lambda(x)$ has compact support. Then W are independent of λ provided that $\lambda(x) = 1$ inside a certain compact region (that depends on the supports of the f_i). This is essentially the cancellation of IR divergences (which come from λ having non-compact support).

Example 15.4. Suppose the union of the supports of the f_i is S, as in the picture.[[★★★ picture]] we need to add little bits x such that it is possible to send a message from x to S and from S to x. ⋄

We need to know that if S is compact, then this larger region is also compact. For Minkowski spacetime, this is an easy exercise, but it is not true for all spacetimes. For example, take $\mathbb{R}^{1,3} \smallsetminus pt$ and make it so that you have to add a neighborhood of the missing point to S. But this spacetime has a hole in it, and you don't know what is coming through the hole, so you don't expect things to work very well. The compactness condition is almost equivalent to spacetime being globally hyperbolic.

So on globally hyperbolic spacetimes (which are the only ones anybody works with), the infrared divergences of quantum field theory automatically cancel.

References

[CdS03] Ana Cannas da Silva. *Introduction to symplectic and Hamiltonian geometry*. Publicações Matemáticas do IMPA. [IMPA Mathematical Publications]. Instituto de Matemática Pura e Aplicada (IMPA), Rio de Janeiro, 2003.

[DEF+99] Pierre Deligne, Pavel Etingof, Daniel S. Freed, Lisa C. Jeffrey, David Kazhdan, John W. Morgan, David R. Morrison, and Edward Witten, editors. *Quantum fields and strings: a course for mathematicians. Vol. 1, 2*. American Mathematical Society, Providence, RI, 1999. Material from the Special Year on Quantum Field Theory held at the Institute for Advanced Study, Princeton, NJ, 1996–1997.